

Locally toric manifolds and singular Bohr-Sommerfeld leaves

Mark Hamilton*
 umbra@math.utoronto.ca

25 September, 2007

Abstract

When geometric quantization is applied to a manifold using a real polarization which is “nice enough”, a result of Śniatycki says that the quantization can be found by counting certain objects, called Bohr-Sommerfeld leaves. Subsequently, several authors have taken this as motivation for counting Bohr-Sommerfeld leaves when studying the quantization of manifolds which are less “nice”.

In this paper, we examine the quantization of compact symplectic manifolds that can locally be modelled by a toric manifold, using a real polarization modelled on fibres of the moment map. We compute the results directly, and obtain a theorem similar to Śniatycki’s, which gives the quantization in terms of counting Bohr-Sommerfeld leaves. However, the count does not include the Bohr-Sommerfeld leaves which are singular. Thus the quantization obtained is different from the quantization obtained using a Kähler polarization.

1 Introduction

Broadly speaking, *quantization* is a procedure which associates to a symplectic manifold M a Hilbert space $\mathcal{Q}(M)$. There are numerous methods of quantization; in this paper we consider *geometric quantization*. The ingredients for geometric quantization are: a symplectic manifold (M, ω) , a complex line bundle \mathbb{L} over M , and a connection ∇ on \mathbb{L} whose curvature is ω . The Hilbert spaces are constructed from sections of \mathbb{L} , or, more generally, higher cohomology groups associated with \mathbb{L} .

*Supported by a PIMS Postdoctoral Fellowship

One additional piece of structure is required, called a *polarization*. This is a method for restricting which sections of \mathbb{L} are considered; it is necessary because the space of all sections is usually “too big.” One example is a Kähler polarization, which is given by a complex structure on M ; the quantization is then determined in terms of holomorphic sections of \mathbb{L} . Another example is a *real polarization*, which is defined more fully below. It is then a natural question to ask, if we have different polarizations on a manifold M , do we obtain the same quantizations from them? This question can be called “independence of polarization.”

A real polarization is given by a foliation of M into Lagrangian submanifolds. The sections of interest are those which are *leafwise flat*: covariant constant (with respect to ∇) in the directions tangent to the leaves of the foliation. If \mathcal{J} is the sheaf of leafwise flat sections, then the quantization is constructed from the cohomology groups $H^k(M; \mathcal{J})$.

If the leaf space B^n is a Hausdorff manifold and the map $\pi: M^{2n} \rightarrow B^n$ is a fibration with compact fibres, a theorem of Śniatycki [S1] says that the above cohomology groups are zero except in dimension n ; furthermore, H^n can be expressed in terms of Bohr-Sommerfeld leaves. A *Bohr-Sommerfeld* leaf is one on which is defined a global section which is flat along the leaf. The set of Bohr-Sommerfeld leaves is discrete, and Śniatycki’s theorem says that the dimension of H^n is equal to the number of Bohr-Sommerfeld leaves. (A similar result holds if the fibres are not compact, except the nonzero group occurs in a different dimension, equal to the rank of the fundamental group of the fibres.)

Quantization using real polarizations, and the relation to the Bohr-Sommerfeld leaves, has also been studied by Andersen [A]. He uses a different approach than we do in this paper, looking at the index theory of a “leafwise complex.” He generalizes Śniatycki’s theorem to regular polarizations which are not necessarily real or Kähler, but something in between.

In many examples of interest, however, what seems like a real polarization is not quite a fibration, but can be viewed as a real polarization with singularities. Several authors, motivated by Śniatycki’s theorem, have defined the quantization in such cases to be that obtained by simply counting Bohr-Sommerfeld leaves. The characterization of Bohr-Sommerfeld leaves often includes fibres which are singular; common practice has been to include the singular fibres in the count, since in many cases this gives a result which agrees with the quantization obtained using a Kähler polarization. (Two examples are Guillemin-Sternberg studying the Gelfand-Cetlin system in [GS1], and Jeffrey-Weitsman studying the moduli space of flat $SU(2)$ connections on a 2-manifold in [JW].)

One example of a “singular real polarization” is the level sets of the moment map on a toric manifold. In this paper, we calculate directly the sheaf cohomology of a toric manifold, or more generally a manifold equipped with a Lagrangian fibration with elliptic singularities, with respect to this singular polarization. The result we obtain is similar to Śniatycki’s theorem: all cohomology groups are zero except in one dimension, and the nonzero group can be computed by counting Bohr-Sommerfeld fibres. However, the singular Bohr-Sommerfeld fibres are *not* included in this count.

Main Theorem (8.10). Let M be a compact symplectic $2n$ -manifold equipped with a locally toric singular Lagrangian fibration, with prequantization line bundle (\mathbb{L}, ∇) . Let \mathcal{J} be the sheaf of leafwise flat sections of \mathbb{L} . Then the cohomology groups $H^k(M; \mathcal{J})$ are zero for all $k \neq n$, and

$$H^n(M; \mathcal{J}) \cong \bigoplus_{b \in BS} \mathbb{C}$$

where the sum is taken over all nonsingular Bohr-Sommerfeld fibres.

This is a surprising result, and is contrary to expectations. It implies that this quantization is different from that computed using a Kähler polarization. For a toric manifold foliated by fibres of the moment map, the Bohr-Sommerfeld leaves correspond to the integer lattice points in the moment polytope. The quantization coming from a Kähler polarization has dimension equal to the number of lattice points in the polytope, including the ones on the boundary. The quantization computed using the methods of this paper has dimension equal to the number of lattice points in the *interior* of the polytope, i.e. excluding the ones on the boundary.

The key point in the calculations where this difference appears is Theorem 4.6, where we calculate the sheaf cohomology of a small neighbourhood of a singularity and find that it is zero. Thus, the singular leaf does not make a contribution to the cohomology, even though it is in some sense a Bohr-Sommerfeld leaf. This calculation relies on the observation in Proposition 4.3 that there are no nonzero leafwise flat sections defined on a neighbourhood of the origin. The underlying reason for this lack of contribution, however, is unclear.

1.1 Methods

The method of proof we use is to compare the manifold to a standard model space. We prove results for the model space by hand, and then apply them locally to the manifold. The types of spaces we consider are compact manifolds

which, roughly, locally look like toric manifolds foliated by leaves of a moment map. (See Definition 2.8 for the precise, if technical, statement.) These include not only toric manifolds, but also integrable systems with elliptic singularities. We consider two types of model spaces: the cylinder $\mathbb{R} \times S^1$, and the complex plane \mathbb{C} .

We equip each model space with a standard prequantization line bundle, and calculate the sheaf cohomology of each by hand (in sections 3 and 4). We show how our methods apply to a simple example, that of S^2 acted on by the circle, in section 5. After defining the model space for higher dimensions in section 6, in section 7 we use a sheaf theoretic argument to patch together the lower-dimensional results. The hypothesis that our manifold possesses a Lagrangian fibration with elliptic singularities guarantees that a neighbourhood of a Bohr-Sommerfeld leaf “looks like” an open set in this model space, in a way that is compatible with the calculation of sheaf cohomology. In this way, we apply (in section 8) the results obtained for the model space to obtain the results for the general manifold. Finally, we return to the specific case of toric manifolds in section 9 to discuss the comparison between real and Kähler polarizations.

Acknowledgements: This paper is based on the author’s Ph.D. thesis, carried out under the supervision of Yael Karshon and Lisa Jeffrey. I am very grateful for all of their support over the course of my Ph.D.

I am grateful to Ruxandra Moraru for suggesting the methods of Chapter 7 and explaining some of the relevant sheaf theory. I also wish to thank Alejandro Uribe for helpful comments; Eva Miranda, Paul Selick, and Pramath Sastry for helpful suggestions and references; and Megumi Harada for much helpful advice.

2 Background

2.1 Connections

Let V be a vector bundle over a manifold M , $\Gamma(V)$ be the space of smooth¹ sections of V , and $\Omega^k(M)$ the space of (smooth) differential k -forms on M .

Definition. Formally, a *connection* on a vector bundle V is a map $\nabla: \Gamma(V) \rightarrow \Omega^1(M) \otimes \Gamma(V)$ which satisfies the following properties:

1. $\nabla(\sigma_1 + \sigma_2) = \nabla\sigma_1 + \nabla\sigma_2$

¹ In this paper, we always take sections, functions, and differential forms to be smooth.

$$2. \quad \nabla(f\sigma_1) = (df) \otimes \sigma_1 + f\nabla\sigma_1$$

for all sections σ_1 and σ_2 and functions f . We typically write $\nabla_X\sigma$ for $\nabla\sigma$ applied to the vector (field) X . This is also called the *covariant derivative* of σ in the direction X .

In order to express a connection in terms which are useful for calculations, we work locally. The following description, taken from [W] (Appendix A.3), only applies to (complex) line bundles, but that is all we require for this paper.

Let \mathbb{L} be a complex line bundle over some manifold M , with s the unit section in some local trivialization over $U \subset M$. Fix a connection ∇ on \mathbb{L} . Define the *potential one-form* Θ of the connection, which is a one-form² on U , by

$$\nabla_X s = -i\Theta(X)s. \quad (2.1)$$

The form Θ gives a complete description of the connection, as follows: any other section σ can be written as $\sigma = fs$ for some complex-valued function f , and so then using (2.1), we obtain

$$\nabla_X \sigma = df(X)s - fi\Theta(X)s. \quad (2.2)$$

Usually, the trivializing section will be implicit, and we will blur the distinction between a section and the complex function representing it.

Note: This description of a connection is only valid over an open set over which the bundle \mathbb{L} is trivializable.

Definition. The *curvature* of the connection is the 2-form Ω on M defined by $\Omega = d\Theta$. This is well-defined, independently of the choice of trivializing section (see Prop. 2.2). A connection is *flat* if its curvature is zero.

The description of a connection in terms of its potential one-form depends on the choice of trivialization. In the following, we compute the effect that changing the trivialization has on the potential one-form.

Proposition 2.1. Let \mathbb{L} be a line bundle with connection ∇ over M . Suppose we have two local trivializations of \mathbb{L} over some open set, with unit sections s and s' , respectively, which are related by $s = \psi s'$. (Here ψ will be a nonvanishing \mathbb{C} -valued function.³) If Θ and Θ' are the potential one-forms with respect to these trivializations, then

$$\Theta' = \Theta - i\frac{1}{\psi}d\psi. \quad (2.3)$$

²If \mathbb{L} is a Hermitian line bundle, and if ∇ and s respect the Hermitian structure, then Θ will in fact be real-valued.

³If the bundles are Hermitian, and the trivializations respect the Hermitian structure, ψ takes values in S^1 .

Proof. Let X be a vector. Then, by eq. (2.1),

$$\nabla_X s = -i\Theta(X)\psi s';$$

also, we have

$$\nabla_X \psi s' = d\psi(X)s' - i\Theta'(X)\psi s'.$$

Equating these and solving gives

$$-i\Theta(X)\psi s' = d\psi(X)s' - i\Theta'(X)\psi s';$$

cancelling common factors (including s') and dividing by ψ (which is never zero) gives

$$\Theta' = \Theta - i\frac{1}{\psi} d\psi \quad (2.4)$$

as one-forms.⁴

□

Proposition 2.2. The curvature form Ω of a connection is independent of the trivialization used to compute the potential one-form.

Proof. This is a straightforward calculation. If we have two different trivializations, the potential one-forms are related by

$$\Theta' = \Theta - i\frac{1}{\psi} d\psi.$$

Thus

$$\begin{aligned} \Omega' &= d\Theta' = d\left(\Theta - i\frac{1}{\psi} d\psi\right) \\ &= d\Theta - i d\left(\frac{1}{\psi}\right) \wedge d\psi + i\frac{1}{\psi} d d\psi \\ &= d\Theta + i\frac{1}{\psi^2} d\psi \wedge d\psi \\ &= d\Theta. \end{aligned} \quad (2.5)$$

This also implies that Ω is defined on all of M .

□

⁴ If we are in the Hermitian case, as in the previous footnotes, then $i\frac{1}{\psi} d\psi$ will be real-valued: since ψ is S^1 -valued, it can be written locally as e^{if} for some real-valued function f , so

$$d\psi = e^{if} idf.$$

Then

$$i\frac{1}{\psi} d\psi = -df$$

is real-valued.

2.1.1 Holonomy

Suppose we have a line bundle with connection (\mathbb{L}, ∇) .

Let γ be a curve on M , with tangent vector $\dot{\gamma}$, and suppose σ is a section of \mathbb{L} defined at least on γ . Let $\tilde{\gamma}$ be the lifting of γ to \mathbb{L} via σ , i.e., $\tilde{\gamma} = \sigma \circ \gamma$. Then the lifting $\tilde{\gamma}$ is *horizontal* if

$$\nabla_{\dot{\gamma}} \sigma = 0 \quad (2.6)$$

for all points along the curve.

Given a curve γ in M and a point x in the fibre over $\gamma(0)$, the curve $\tilde{\gamma}$ is uniquely determined by the condition that it is a horizontal lift of γ with $\tilde{\gamma}(0) = x$. Taking x to $\tilde{\gamma}(1)$ gives a map from $\mathbb{L}_{\gamma(0)}$ to $\mathbb{L}_{\gamma(1)}$, called *parallel transport* along γ . If γ is a loop, this gives an automorphism of $\mathbb{L}_{\gamma(0)}$, called the *holonomy* around γ .

If \mathbb{L} is Hermitian and the connection respects the Hermitian structure, we can view holonomy as a map from $\{\text{loops on } M\}$ to S^1 . This map is given by

$$\text{hol}: \gamma \mapsto \exp \left(i \int_{\gamma} \Theta \right). \quad (2.7)$$

If the connection is flat, then homotopic loops have the same holonomy (the proof is basically Stokes' theorem) and the holonomy can be viewed as a map from $\pi_1(M)$ to the automorphisms of the fibre.

2.2 Sheaves and cohomology

We review the definition of sheaves and the construction of Čech cohomology, as they will be used extensively in this paper. The material in this section is standard, and so we do not prove our assertions. See, for example, [GH], pp. 38–40.

Roughly speaking, a sheaf is a collection of functions on open sets, often satisfying some further conditions (holomorphic, etc.). The precise definition is as follows:

Definition. Let X be a topological space. A presheaf (of abelian groups) \mathcal{F} on X assigns to every open set U of X an abelian group $\mathcal{F}(U)$, usually referred to as the set of sections of \mathcal{F} over U . It also assigns restriction maps: to any $V \subset U$, the presheaf assigns a map $\mathcal{F}(U) \rightarrow \mathcal{F}(V)$, “restriction to V ”, such that if $W \subset V \subset U$ and $\sigma \in \mathcal{F}(U)$, then

$$\sigma|_W = (\sigma|_V)|_W,$$

and if $V = U$ then “restriction” is just the identity map.

Definition. A presheaf \mathcal{S} is a *sheaf* if the following properties hold:

1. For any pair of open sets U, V , and sections $\sigma \in \mathcal{S}(U)$ and $\tau \in \mathcal{S}(V)$ which agree on the intersection $U \cap V$, there exists a section $\rho \in \mathcal{S}(U \cup V)$ which restricts to σ on U and τ on V .
2. If σ and τ in $\mathcal{S}(U \cup V)$ have equal restrictions to U and V , then they are equal on $U \cup V$.

Definition 2.3. For a sheaf \mathcal{S} over M , the *stalk* of \mathcal{S} over $x \in M$ is

$$\mathcal{S}_x = \varinjlim \mathcal{S}(U)$$

where the limit is taken over all open sets U containing x .

The Čech cohomology of M with coefficients in the sheaf \mathcal{S} (or just “the sheaf cohomology of M ”)⁵ is defined as follows.

Fix an open cover $\mathcal{A} = \{A_\alpha\}$ of M .

A Čech k -cochain assigns, to each $(k + 1)$ -fold intersection of elements from the cover \mathcal{A} , a section of \mathcal{S} . We write $A_{\alpha_0 \dots \alpha_k}$ for $A_{\alpha_0} \cap \dots \cap A_{\alpha_k}$, where the α_j are distinct. Then a k -cochain is an assignment $f_{\alpha_0 \dots \alpha_k} \in \mathcal{S}(A_{\alpha_0 \dots \alpha_k})$ for each $(k + 1)$ -fold intersection in the cover \mathcal{A} . Denote the set of k -cochains by $C_{\mathcal{A}}^k(M; \mathcal{S})$, or just $C_{\mathcal{A}}^k$ when the manifold and sheaf are understood.

Next, we define a coboundary operator δ to make $C_{\mathcal{A}}^*$ into a cochain complex. For $f = \{f_{\alpha_0 \dots \alpha_{k-1}}\}$ a $(k - 1)$ -cochain, δf will be a k -cochain; thus, we need to give a section corresponding to each $(k + 1)$ -fold intersection. This is done as follows:

$$(\delta f)_{\alpha_0 \dots \alpha_k} = \sum_{j=0}^k (-1)^j f_{\alpha_0 \dots \hat{\alpha}_j \dots \alpha_k}|_{A_{\alpha_0 \dots \dots \alpha_k}} \quad (2.8)$$

where the $\hat{\alpha}$ denotes that the index is omitted. Thus, for instance, $(\delta f)_{123} = f_{23} - f_{13} + f_{12}$, all restricted to A_{123} .

A (straightforward, but tedious) calculation shows that $\delta \circ \delta = 0$, and so $C_{\mathcal{A}}^*$ is a cochain complex. The sheaf cohomology with respect to the cover \mathcal{A} is the cohomology of this complex,

$$H_{\mathcal{A}}^k(M; \mathcal{S}) = \frac{\ker \delta^k}{\text{im } \delta^{k-1}}$$

⁵ The “sheaf cohomology of M ” is actually defined in a more abstract way using resolutions of the sheaf \mathcal{S} . However, for manifolds, the result obtained is the same as the Čech cohomology, and the latter is more convenient for our calculations.

(where by δ^k we mean the map δ on $C_{\mathcal{A}}^k$).

Another cover \mathcal{B} is a *refinement* of \mathcal{A} , and we write $\mathcal{B} \leq \mathcal{A}$, if every element of \mathcal{B} is a subset of some element of \mathcal{A} . From this we define a map $\phi: C_{\mathcal{A}}^k(U, \mathcal{S}) \rightarrow C_{\mathcal{B}}^k(U, \mathcal{S})$, induced by the restriction maps in the sheaf: simply restrict each element of a cochain, defined on some intersection of sets in \mathcal{A} , to the intersection of the corresponding sets in \mathcal{B} . More formally, a refinement gives a map $\rho: \mathcal{B} \rightarrow \mathcal{A}$, where $B \subset \rho(B)$ for all $B \in \mathcal{B}$. Then, if $\eta \in C_{\mathcal{A}}^k$ is a cochain, $\phi\eta$ is defined by

$$(\phi\eta)_{B_0 B_1 \dots B_k} = (\eta)_{(\rho B_0)(\rho B_1) \dots (\rho B_k)}|_{B_0 B_1 \dots B_k}.$$

Since this map is essentially just restriction, it commutes with δ , and so it induces a map on cohomology $H_{\mathcal{A}}^* \rightarrow H_{\mathcal{B}}^*$. Two different choices of maps ρ for the same refinement induce chain homotopic maps on cochains, and thus induce the same map on cohomology. These maps turn the collection of $H_{\mathcal{A}}^*$ for all open covers of M into a directed system.

Finally, the honest-to-goodness *sheaf cohomology* of M is defined as the limit of this directed system:

$$H^*(M; \mathcal{S}) = \varinjlim H_{\mathcal{A}}^*(M; \mathcal{S}).$$

Lemma 2.4. Suppose \mathfrak{B} is a collection of covers \mathcal{B} of M such that any open cover \mathcal{A} of M has a refinement $\mathcal{B} \in \mathfrak{B}$. Suppose furthermore that all $H_{\mathcal{B}}^*(M; \mathcal{S})$ are isomorphic, for each $\mathcal{B} \in \mathfrak{B}$. Then the sheaf cohomology of M is isomorphic to the cohomology computed using one of the covers \mathcal{B} .

Proof. This follows easily from the construction of the direct limit. The set \mathfrak{B} is *cofinal* in the set of all covers if every cover has a refinement from \mathfrak{B} . \square

2.2.1 Naturality of Direct Limits and Cohomology

The following results will be used in Section 3.4.

Lemma 2.5 ([G], p. 10). The direct limit of exact sequences is exact. More precisely, suppose we have:

- three directed systems of modules L'_i , L_i , and L''_i (indexed by the same directed set I), with direct limits L' , L , and L'' respectively, and
- for each i , an exact sequence $L'_i \rightarrow L_i \rightarrow L''_i$.

Suppose further that, for each $i \geq j$, the diagram

$$\begin{array}{ccccc} L'_i & \longrightarrow & L_i & \longrightarrow & L''_i \\ \downarrow & & \downarrow & & \downarrow \\ L'_j & \longrightarrow & L_j & \longrightarrow & L''_j \end{array}$$

where the vertical arrows denote the maps in the directed system, is commutative.

Then the sequence $L' \rightarrow L \rightarrow L''$ is exact.

Lemma 2.6 ([MacL], Prop. 4.2). Given a transformation of short exact sequences of cochain complexes

$$\begin{array}{ccccccc} 0 & \longrightarrow & A^* & \longrightarrow & B^* & \longrightarrow & C^* \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & D^* & \longrightarrow & E^* & \longrightarrow & F^* \longrightarrow 0 \end{array}$$

this induces a map between the long exact sequences of cohomology

$$\begin{array}{ccccccc} \cdots & \longrightarrow & H^*(A) & \longrightarrow & H^*(B) & \longrightarrow & H^*(C) \longrightarrow \cdots \\ & & \downarrow & & \downarrow & & \downarrow \\ \cdots & \longrightarrow & H^*(D) & \longrightarrow & H^*(E) & \longrightarrow & H^*(F) \longrightarrow \cdots \end{array}$$

2.3 Toric manifolds

We briefly review the definition of, and a few facts about, toric manifolds, referring the reader to [CdS1] for a more detailed introduction.

Definition. A *toric manifold* is a compact symplectic manifold M of dimension $2n$ equipped with an effective Hamiltonian action of the torus $T^n = (S^1)^n$.

Recall that if a symplectic manifold M^{2n} has an effective Hamiltonian action of a torus T^k , then $k \leq n$. Thus M is toric if T has maximal dimension.

The “Hamiltonian” in the definition means that a toric manifold is equipped with a T -invariant *moment map* $\mu: M \rightarrow \mathbb{R}^n \cong \mathfrak{t}^*$, which satisfies the following condition: For $\xi \in \mathfrak{t}$, let ξ_M be the generating vector field for the action on M , and let $\mu^\xi: M \rightarrow \mathbb{R}$ be the “component of μ in the ξ direction,” that is,

$$\langle \mu(x), \xi \rangle = \mu^\xi(x)$$

for all $x \in M$, where $\langle \cdot, \cdot \rangle$ denotes the pairing between \mathfrak{t} and \mathfrak{t}^* . Then

$$\iota_{\xi_M} \omega = d\mu^\xi, \quad (2.9)$$

that is, ξ_M is the Hamiltonian vector field of μ^ξ . We do not use this definition directly, but some facts about toric manifolds.

First, fibres of the moment map are orbits of the torus action. Second, a theorem of Atiyah and Guillemin-Sternberg says that the image of the moment map is a convex polytope Δ in $\mathfrak{t}^* \cong \mathbb{R}^n$. If $x \in \Delta$ lies on a face of codimension k , then the stabilizer of a point in $\mu^{-1}(x)$ is a torus of dimension k , and $\mu^{-1}(x)$ is a torus of dimension $n - k$. (This is Lemma 2.2 in [D].) For x in the interior of Δ , the fibres are n -dimensional tori; the moment map condition implies they are Lagrangian.

The following lemma, known as the Local Normal Form, gives a description of the neighbourhood of an orbit.

Lemma 2.7 (Local Normal Form, Lemma 2.5 in [D]). Let M be a symplectic toric manifold with moment polytope Δ and moment map μ . Let F be a face of Δ of dimension m ; V a convex, open, relatively compact subset of F ; x a point in V ; and D a ball around 0 in $\mathbb{C}^{2(n-m)}$. Let ω_0 be the following symplectic form on $(S^1)^m \times V \times D$:

$$\omega_0 = \sum_{1 \leq j \leq m} d\alpha_j \wedge da_j + \sum_{m+1 \leq j \leq n} dx_j \wedge dy_j$$

where α are coordinates on S^1 , a are coordinates on V , and $z = x + iy$ are coordinates on D .

Then there is a symplectic isomorphism from a neighbourhood of $\mu^{-1}(V)$ onto $(S^1)^m \times V \times D$ taking the action of T to the action of $(S^1)^n$ defined by

$$\begin{aligned} (\theta_1, \dots, \theta_n) \cdot (\alpha_1, \dots, \alpha_m, a_1, \dots, a_m, z_{m+1}, \dots, z_n) = \\ (\alpha_1 + \theta_1, \dots, \alpha_m + \theta_m, a_1, \dots, a_m, e^{i\theta_{m+1}} z_{m+1}, \dots, e^{i\theta_n} z_n) \end{aligned} \quad (2.10)$$

with moment map

$$\mu = p + (a_1, \dots, a_m, |z_{m+1}|^2, \dots, |z_n|^2).$$

2.4 Geometric quantization and polarizations

Let (M, ω) be a compact symplectic manifold of dimension $2n$.

2.4.1 Quantization

The theory of geometric quantization was initiated by Kostant and Souriau in the 1970s, and remains an active area of research today, with applications to both physics and representation theory. We do not attempt to give a comprehensive introduction here. For a nice, brief overview of the ideas behind it, see [GGK], section 6.1, or [GS2], chapter 34. For a more thorough introduction, see [P]. Two classic references, albeit somewhat technical, are [S2] and [W].

The basic idea of quantization is to associate to a symplectic manifold (M, ω) a Hilbert space (or a vector space) $\mathcal{Q}(M)$. (The terminology ‘quantization’ comes from physics, where we think of M as a classical mechanical system, and $\mathcal{Q}(M)$ as the space of wave functions of the corresponding quantum system.) Much of the motivation for geometric quantization in mathematics comes from representation theory.

In geometric quantization, the quantum space is constructed from the sections of a complex line bundle or, more generally, from higher-dimensional cohomology groups associated with the line bundle.

Definition. A *prequantization line bundle* over M is a Hermitian line bundle \mathbb{L} over M , with a connection ∇ whose curvature is ω . M is *prequantizable* if it possesses a prequantization line bundle. (This will be the case iff the symplectic form satisfies the integrality condition that $\frac{1}{2\pi}[\omega] \in H^2(m, \mathbb{Z})$. See, for example, [W], section 8.3.⁶)

We would like $\mathcal{Q}(M)$ to be the space of sections of \mathbb{L} . However, this space is generally ‘too big.’ As noted in the Introduction, the solution is to use a ‘polarization’ to choose a subspace of the space of sections; the quantum space is then constructed from only ‘polarized’ sections.

Our main interest in this paper is in *real polarizations*.⁷ The usual definition of a real polarization on M is a sub-bundle $P \subset TM$ which is Lagrangian and integrable. In our case, we allow ‘singular polarizations’, where the leaves are not all of the same dimension. We define a locally toric singular Lagrangian fibration to be a structure which locally looks like the (singular) fibration of a toric manifold by the moment map, using the local structure given by Lemma 2.7.

⁶ The exact form of the integrality condition depends on the conventions used, and one may see instead $\frac{1}{\hbar}[\omega]$, $\frac{1}{2\pi\hbar}[\omega]$, or just $[\omega]$ required to be integral. We take $\hbar = 1$, and use coordinates on S^1 running from 0 to 2π , which gives the integrality condition stated here.

⁷ Thus in particular we do not give the general theory of polarizations, but refer the reader to [S2], pages 8–11, or [W] sections 4.5 and 5.4. We will also have occasion to mention Kähler polarizations, which we define below.

Definition 2.8. A *locally toric singular Lagrangian fibration* on a symplectic $2n$ -manifold M is a map $\pi: M \rightarrow B$ to a topological space B such that for every point in B , there exist:

- a nonnegative integer k
- a neighbourhood $U \subset B$ of the point
- an open subset $\Omega \subset \mathbb{R}^{n-k} \times \mathbb{R}_+^k$
- a homeomorphism $\psi: U \cong \Omega$
- a symplectomorphism $\tilde{\psi}: \pi^{-1}(U) \cong \pi_0^{-1}(\Omega)$

such that the following diagram commutes:

$$\begin{array}{ccc} \pi^{-1}(U) & \xrightarrow{\tilde{\psi}} & (S^1 \times \mathbb{R})^{n-k} \times \mathbb{C}^k \\ \pi \downarrow & & \pi_0 \downarrow \\ U & \xrightarrow{\psi} & \mathbb{R}^{n-k} \times \mathbb{R}_+^k \end{array} \quad (2.11)$$

where $\pi_0: (S^1 \times \mathbb{R})^{n-k} \times \mathbb{C}^k \rightarrow \mathbb{R}^{n-k} \times \mathbb{R}_+^k$ is the projection to \mathbb{R} on the first $n-k$ factors, and the projection $(x, y) \mapsto \frac{1}{2}(x^2 + y^2)$ in the last k factors, and where we take the standard symplectic structure on \mathbb{C} and $S^1 \times \mathbb{R} \cong T^*S^1$.

Remark. This implies that on an open subset of B , the preimages of points are Lagrangian manifolds. The level sets of π form a singular Lagrangian fibration. The singular fibres are those with $k > 0$. Note, however, that the fibres are only “singular” in terms of the fibration. They are still smooth manifolds.

Definition. A *locally toric singular real polarization* on M is the distribution $P \subset TM$ associated to a locally toric singular Lagrangian fibration (i.e., P_x is the set of directions tangent to the leaf through x).

This definition includes toric manifolds, by the Local Normal Form. In this case the map π is the moment map, and B can be taken to be the moment polytope.

It also includes more general integrable systems with certain kinds of singularities. Eliasson in [E1] and [E2] and Miranda in [Mr] established the local symplectic classification of non-degenerate singularities of integrable Hamiltonian systems: such singularities are isomorphic to the product of singularities of three basic types, called *elliptic*, *hyperbolic*, and *focus-focus*. Definition 2.8

includes integrable systems which have only elliptic singularities, by the following theorem (due to Dufour and Molino and Eliasson, here taken from Zung [Z]):

Theorem 2.9 (3.9 in [Z]). Let N be an elliptic singular leaf of codimension k in an integrable system with moment map $F: M^{2n} \rightarrow \mathbb{R}^n$. Then on a tubular neighbourhood of N , there exist symplectic coordinates $(x_1, \dots, x_n, y_1, \dots, y_n)$ so that:

- y_1, \dots, y_{n-k} are mod 1
- $\omega = \sum dx_j \wedge dy_j$
- $N = \{x_1 = \dots = x_n = y_{n-k+1} = \dots = y_n = 0\}$, i.e. N is a $n - k$ -torus, with coordinates y_1, \dots, y_{n-k}
- F is a smooth function of x_1, \dots, x_{n-k} , and $x_j^2 + y_j^2$ for $n - k < j \leq n$.

(The relation of the notation in this theorem to that used in this paper is as follows: k represents the same thing. What we call t_j, θ_j in Section 3 correspond to x_j, y_j (except for factors of 2π) for $j \leq n - k$. What we call s_j in Section 6 is $\frac{1}{2}(x_i^2 + y_i^2)$ in the coordinates in this theorem, for $1 \leq j \leq k$, $i = (n - k) + j$.)

Kogan in [Kg] gives a description of the structure of these spaces.

Definition 2.10. Given a manifold M with prequantization line bundle \mathbb{L} and (possibly singular) real polarization P , a section σ of \mathbb{L} over $U \subset M$ is *flat along the leaves*, or *leafwise flat*, if $\nabla_X \sigma = 0$ for all $X \in P$, at every point of U .

We denote the sheaf of leafwise flat sections on M by \mathcal{J}_M (or just \mathcal{J}).

Definition. The *quantization* of M is the sum of the cohomology groups of M with coefficients in the sheaf of leafwise flat sections:

$$\mathcal{Q}(M) := \bigoplus_q H^q(M; \mathcal{J}) \tag{2.12}$$

Thus, in this paper the central items of interest are the sheaf cohomology spaces $H^q(M; \mathcal{J})$.

Remark. Various authors define the quantization in terms of sheaf cohomology, either as the direct sum as in (2.12) (e.g. in [JW]), or as the alternating sum of cohomology (e.g. in [GGK]). In [GS1] the authors call the groups $H^q(M; \mathcal{J})$ “the basic quantum data associated with $M\dots$,” without defining

the quantization as either sum. However, in all of these cases, as well as in this paper, all but one of the groups turn out to be zero, and so whether one takes the direct or the alternating sum doesn't matter in the end. We use the convention of (2.12), and call the resulting object "the quantization of M ."

2.4.2 Bohr-Sommerfeld leaves and Śniatycki's theorem

Let M be a prequantizable compact symplectic manifold of dimension $2n$, as in the previous section.

Definition. A leaf ℓ of the polarization P is a *Bohr-Sommerfeld leaf* if there exists a globally defined nonzero section of \mathbb{L} along ℓ , whose covariant derivative (with respect to ∇) is zero in directions tangent to P . (Here "globally defined" means defined on all of ℓ , not all of M .)

The *Bohr-Sommerfeld set* is the set of points in B whose preimages are Bohr-Sommerfeld leaves.

In [S1], Śniatycki proves that, in the case where the projection map $\pi: M \rightarrow B$ is a fibration, the cohomology groups $H^q(M; \mathcal{J})$ appearing in (2.12) are all zero except in dimension n . Furthermore, H^n can be computed by counting Bohr-Sommerfeld leaves. More precisely, we have the following result:

Theorem (Śniatycki, 1975 [S1]). Let M be a $2n$ -dimensional symplectic manifold, with a prequantization line bundle \mathbb{L} as above. Let P be a real polarization such that the projection map $\pi: M \rightarrow B$ is a fibration with compact fibres. Then $H^q(M; \mathcal{J}) = 0$ for all $q \neq n$.

Furthermore, let $\Gamma_{BS}(\mathbb{L})$ be the space of smooth sections of \mathbb{L} along the union of Bohr-Sommerfeld leaves, and $C_P^\infty(M)$ be the ring of functions on M constant on leaves of π . Then provided P satisfies an orientability condition, $H^n(M; \mathcal{J})$ is isomorphic to $\Gamma_{BS}(\mathbb{L})$, as modules over $C_P^\infty(M)$.

More generally, if the leaves are not compact, then similar results are true with n replaced by the rank of the fundamental group of a typical integral manifold of P .

As a vector space, $H^n(M; \mathcal{J})$ is isomorphic to the direct sum of copies of \mathbb{C} , with one copy for each Bohr-Sommerfeld leaf.

2.5 Examples

Example. A simple example of a compact space to which Śniatycki's results apply is T^4 fibred over T^2 . If T^4 has coordinates (x_1, x_2, y_1, y_2) , all mod 2π ,

and standard symplectic form, then projecting to T^2 via

$$(x_1, x_2, y_1, y_2) \mapsto (x_1, x_2)$$

is a Lagrangian fibration.

Example. A less trivial example is Thurston's example [Th] of a symplectic manifold M_Θ which is not Kähler (and therefore not toric). Cannas da Silva in [CdS2] gives the following description of M_Θ :

Let Γ be the discrete group generated by the following symplectomorphisms of \mathbb{R}^4 :

$$\begin{aligned}\gamma_1 &= (x_1, x_2, y_1, y_2) \mapsto (x_1, x_2 + 1, y_1, y_2) \\ \gamma_2 &= (x_1, x_2, y_1, y_2) \mapsto (x_1, x_2, y_1, y_2 + 1) \\ \gamma_3 &= (x_1, x_2, y_1, y_2) \mapsto (x_1 + 1, x_2, y_1, y_2) \\ \gamma_4 &= (x_1, x_2, y_1, y_2) \mapsto (x_1, x_2 + y_2, y_1 + 1, y_2)\end{aligned}$$

Then $M_\Theta = \mathbb{R}^4/\Gamma$, with symplectic form $\omega = dx_1 \wedge dy_1 + dx_2 \wedge dy_2$.

If we map M_Θ to the 2-torus T^2 by

$$(x_1, x_2, y_1, y_2) \mapsto (y_1, y_2)$$

(all coordinates taken mod 1) then the fibres are tori in the (x_1, x_2) coordinates, which are thus in fact Lagrangian submanifolds, so M_Θ is a 2-torus bundle over the 2-torus. This fibration is a non-singular Lagrangian foliation, and so M_Θ is a compact manifold that satisfies the hypotheses of Śniatycki's theorem.

Example. We can use Thurston's manifold to construct further (admittedly somewhat artificial) examples of manifolds which satisfy Definition 2.8 but are not toric: simply take the product $M_T \times M_\Theta$ of Thurston's manifold with any compact toric manifold M_T . These do not have a global torus action of maximal dimension, because of the M_Θ factor, but they are still locally toric, and so our results apply. Śniatycki's theorem does not apply, however, since the foliation is singular.

2.6 Aside: Rigidity of Bohr-Sommerfeld leaves

Definition. We say that M satisfies *Bohr-Sommerfeld rigidity* if the Bohr-Sommerfeld leaves of M are independent of the choice of prequantum connection on \mathbb{L} .

Proposition 2.11. Let $i: \ell \hookrightarrow M$ be the inclusion of a leaf of the polarization. If the induced map $i_*: H_1(\ell, \mathbb{Z}) \rightarrow H_1(M, \mathbb{Z})$ is zero for all leaves ℓ , then M satisfies Bohr-Sommerfeld rigidity.

Proof. If i_* is zero, this means that any loop γ on ℓ is homotopic to a point in M .

Let Σ be a surface spanning γ . For definiteness, assume that the prequantization connection has a potential one-form Θ defined on all of Σ . Then from (2.7), the holonomy around γ is given by

$$\text{hol}_\gamma = \exp \left(i \int_\gamma \Theta \right)$$

which equals

$$\exp \left(i \int_\Sigma \omega \right) \tag{2.13}$$

by Stokes' theorem. (If there is no Θ defined on all of Σ , break Σ up into little surfaces over which Θ exists, apply Stokes' on each one, and piece back together.)

Now ℓ is a Bohr-Sommerfeld leaf if and only if $\text{hol}_\gamma = 1$ for all loops γ on ℓ . By (2.13), this will be true if and only if

$$\exp \left(i \int_\Sigma \omega \right)$$

is 1 for all Σ which span a loop on ℓ . This depends only on ω and ℓ , not on the connection form. \square

Corollary 2.12. A compact symplectic toric manifold has rigid Bohr-Sommerfeld leaves.

Proof. All odd-degree homology groups of a compact symplectic toric manifold are zero. (See, for example, Theorem I.3.6 in [CdS1].) Thus the image of $H_1(\ell, \mathbb{Z})$ in $H_1(M, \mathbb{Z})$ is certainly zero, and so M satisfies Bohr-Sommerfeld rigidity. \square

Remark. Note that $\mathbb{R} \times S^1$, considered in the next section, does not satisfy rigidity: if we change the connection by adding a (non-integer) constant multiple of $d\theta$ to it, this changes the Bohr-Sommerfeld leaves.

3 The cylinder

The first model space we consider is the cylinder $\mathbb{R} \times S^1$. In this section we compute its sheaf cohomology by hand. For this section, let M denote $\mathbb{R} \times S^1$, with coordinates (t, θ) , where θ is taken mod 2π , and symplectic form $\omega = dt \wedge d\theta$.

In 3.1, we give the basic definition and set-up of the manifold, describe the sheaf of sections flat along the leaves, and calculate the Bohr-Sommerfeld leaves. In the next several sections, we find the sheaf cohomology of a simple type of subset: a “band” around the cylinder. Our strategy will be to find the cohomology of simple “bands” directly, and then piece them together using a Mayer-Vietoris sequence.

We start by computing the cohomology with respect to simple kinds of covers (in 3.2 and 3.3). Next, in 3.4, we derive a Mayer-Vietoris sequence for our particular sheaf. Finally, in 3.5, we take the limit over finer and finer covers to find the actual sheaf cohomology. (Cf. section 2.2.)

3.1 Flat sections and Bohr-Sommerfeld leaves

M is equipped with a real polarization, given by vectors tangent to the S^1 directions. The leaves of the polarization are the fibres of the projection $M \rightarrow \mathbb{R}$.

Let $\mathbb{L} = M \times \mathbb{C}$ be the trivial bundle over M . Let σ be a section of \mathbb{L} , which we can view as a complex-valued function. Since $\omega = dt \wedge d\theta = d(t d\theta)$, a connection with potential one-form $t d\theta$ has curvature ω . Thus the connection given by

$$\nabla_X \sigma = X(\sigma) - \sigma i t d\theta(X) \quad (3.1)$$

makes \mathbb{L} into a prequantization line bundle over M . (Note that, in this case, the potential one-form is defined on the entire manifold.)

In order to calculate the sheaf cohomology, we need to know which sections of \mathbb{L} are flat along the leaves (see Definition 2.10; as there, we denote the sheaf of such sections by \mathcal{J}).

Proposition 3.1. Let $U \subset M$, and let a section of \mathbb{L} over U be given by a map $\sigma: U \rightarrow \mathbb{C}$. Then σ is flat along the leaves if and only if it locally has the form

$$\sigma = a(t) e^{it\theta} \quad (3.2)$$

for some smooth function $a(t)$.

Proof. This result follows directly from the description of the connection given above in (3.1). The section σ is flat along the leaves if

$$0 = \nabla_X \sigma = X(\sigma) - \sigma it d\theta(X) \quad \forall X \in P.$$

The polarization P is the span of $\frac{\partial}{\partial \theta}$, so this condition is equivalent to

$$0 = \nabla_{\frac{\partial}{\partial \theta}} \sigma = \frac{\partial \sigma}{\partial \theta} - \sigma it,$$

i.e.

$$\frac{\partial \sigma}{\partial \theta} = it\sigma.$$

This is a differential equation for $\sigma(t, \theta)$, which is easily solved, giving (3.2). \square

Recall from Section 2.4.2 that a leaf ℓ satisfies the *Bohr-Sommerfeld condition* if it possesses a global covariant constant section (other than zero), i.e. a section defined on the whole leaf which is flat along the leaf.

Lemma 3.2. The Bohr-Sommerfeld set is $\mathbb{Z} \subset \mathbb{R}$, and the Bohr-Sommerfeld leaves are $\{m\} \times S^1$, $m \in \mathbb{Z}$.

Proof. Fix a leaf $\ell_0 = \{t_0\} \times S^1$. By (3.2), a flat section σ over ℓ_0 is of the form $ce^{it_0\theta}$ for some constant c . It will always exist on a small neighbourhood in the leaf, but will only be defined on the whole leaf if its values are compatible as it wraps around the leaf. That is, values of σ must agree for values of θ that differ by 2π . This requires $e^{it_0\theta} = e^{it_0(\theta+2\pi)}$, i.e. $e^{2\pi it_0} = 1$, i.e. $t_0 \in \mathbb{Z}$. \square

Note that the space of global covariant constant sections over one leaf is one-dimensional: $\{\sigma = ce^{it_0\theta} \mid c \in \mathbb{C}\}$.

Let $I \subset \mathbb{R}$ be an open interval, and let $U = I \times S^1 \subset M$. By Śniatycki's theorem,

$$H^1(U, \mathcal{J}) \cong \bigoplus_{m \in \mathbb{Z} \cap I} \mathbb{C}; \quad H^k(U, \mathcal{J}) = 0, \quad k \neq 1.$$

In the following sections we compute the sheaf cohomology of U directly, and show that it agrees with Śniatycki's theorem.

3.2 Sheaf cohomology

In this section we compute the Čech cohomology, with coefficients in \mathcal{J} , of a band in M , with respect to a particular cover. As a warm-up, in this section we use the simplest possible cover; however, as we will see later, this case already shows all the important features of the calculation.

Definition. A *band* is a subset of M of the form $I \times S^1$, with $I \subset \mathbb{R}$ a bounded open interval.

Let U be a band around the cylinder that contains at most one Bohr-Sommerfeld leaf. Partition U into three rectangles E, F , and G by partitioning S^1 into three intervals. (See Figure 1, where the heavy line indicates an overlap.) We will calculate the cohomology of U with respect to the cover $\mathcal{E}_3 = \{E, F, G\}$.

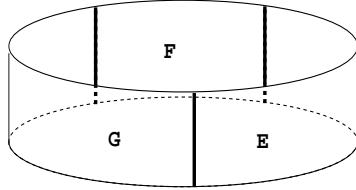


Figure 1: The cover \mathcal{E}_3

3.2.1 H^0

We can see directly that $H^0(U; \mathcal{J}) = 0$, since $H^0(U; \mathcal{J})$ is just the space of global sections of \mathcal{J} over U , which we know from the argument in the proof of Lemma 3.2 is $\{0\}$. However, we will calculate it directly to begin seeing how the Čech approach works in this situation.

A Čech 0-cochain η is an assignment, to each of the sets E, F , and G , of a flat section over that set. Such a section will have the form $a(t)e^{it\theta}$. Index the section by the set, so the piece of η on E is $a_E(t)e^{it\theta}$, etc.

On each set, the coordinate θ can be defined, even though can not be defined on all of S^1 . Fix a branch of θ on each open set W , and denote it by θ_W . Choose these branches so that $\theta_F = \theta_E$ on $E \cap F$, $\theta_G = \theta_F$ on $F \cap G$, and $\theta_G = \theta_E + 2\pi$ on $G \cap E$.

The coboundary of η is the collection

$$(\delta\eta)_{VW} = \eta_W - \eta_V = a_W(t) e^{it\theta_W} - a_V(t) e^{it\theta_V},$$

and so η will be a cocycle if each of these are zero. Applying this to each of the three sets, we have that η is a cocycle iff

$$\begin{aligned} 0 &= a_F(t) e^{it\theta_F} - a_E(t) e^{it\theta_E} && \text{on } E \cap F \\ 0 &= a_G(t) e^{it\theta_G} - a_F(t) e^{it\theta_F} && \text{on } F \cap G \\ 0 &= a_E(t) e^{it\theta_E} - a_G(t) e^{it\theta_G} && \text{on } G \cap E \end{aligned} \tag{3.3}$$

In the first two of these equations, the θ coordinates are equal on the intersections, and so we can cancel the exponential factors; this implies

$$a_E(t) = a_F(t) \quad \text{and} \quad a_F(t) = a_G(t). \quad (3.4)$$

However, θ_E and θ_G differ by 2π on the intersection, and so the third equation becomes

$$0 = a_E(t) e^{it\theta_E} - a_G(t) e^{it\theta_E + 2\pi it},$$

which implies

$$a_E(t) = e^{2\pi it} a_G(t). \quad (3.5)$$

Conditions (3.4) and (3.5) together require that $e^{2\pi it} = 1$, which cannot happen on an interval of t -values. Thus there are no 0-cocycles, and $H^0 = 0$.

3.2.2 H^1

The one-dimensional case is more interesting. A 1-cochain η is an assignment of a flat section η_{VW} to each intersection $V \cap W$; thus a 1-cochain is given by a triple of functions

$$\{a_{EF}(t)e^{it\theta}, a_{FG}(t)e^{it\theta}, a_{GE}(t)e^{it\theta}\}.$$

Since there are no triple intersections in this cover, there are no 2-cochains, and thus every 1-cochain is a cocycle.

Note that, since η is determined by the a 's, and each a is a smooth function of t on I , the space of cocycles is isomorphic to $C^\infty(I)^3$.

We now consider when a 1-cochain is a coboundary, namely, when there exists a 0-cochain $\beta = \{b_E e^{it\theta_E}, b_F e^{it\theta_F}, b_G e^{it\theta_G}\}$ with $\delta\beta = \eta$. This requires that

$$\eta_{VW} = \beta_W - \beta_V \quad \text{on } V \cap W, \quad (3.6)$$

for each pair of V and W .

Notation: We will write EF for $E \cap F$, and so on. The order in which intersections are written matters in Čech cohomology, for bookkeeping, and so we need to be consistent. We will write intersections in the order EF , FG , GE . Furthermore, we will use the convention that, on any intersection VW , we use the θ coordinate from V .

Applying these conventions to the possible equations (3.6), and using the formulas for η and β , we obtain that $\delta\beta = \eta$ iff

$$a_{EF}(t)e^{it\theta_E} = b_F(t)e^{it\theta_F} - b_E(t)e^{it\theta_E} \quad \text{on } EF \quad (3.7a)$$

$$a_{FG}(t)e^{it\theta_F} = b_G(t)e^{it\theta_G} - b_F(t)e^{it\theta_F} \quad \text{on } FG \quad (3.7b)$$

$$a_{GE}(t)e^{it\theta_G} = b_E(t)e^{it\theta_E} - b_G(t)e^{it\theta_G} \quad \text{on } GE \quad (3.7c)$$

In each of these equations, all the θ coordinates are equal on the relevant set, except in (3.7c), where they differ by a factor of 2π . Applying this fact, and cancelling common factors of $e^{it\theta}$, we obtain the following system of three equations

$$\begin{aligned} a_{EF} &= b_F - b_E \\ a_{FG} &= b_G - b_F \\ a_{GE} &= e^{-2\pi it} b_E - b_G \end{aligned} \tag{3.8}$$

in the three unknown functions b_E , b_F , and b_G on I . They must be true for each value of t in I . The matrix of this system is

$$\begin{bmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ e^{-2\pi it} & 0 & -1 \end{bmatrix} \tag{3.9}$$

which is invertible precisely when $e^{-2\pi it} \neq 1$. Thus, by elementary linear algebra, the system (3.8) has a solution, for any collection of the a_{VW} , if $e^{-2\pi it}$ is never 1 on U . In this case, every cocycle is a coboundary, and U has trivial cohomology.

The other possibility is if $e^{2\pi it} = 1$ somewhere in I , which only occurs if I contains an integer m . In this case, by further linear algebra, the system (3.8) only has a solution if η satisfies the condition

$$a_{EF}(m) + a_{FG}(m) + a_{GE}(m) = 0. \tag{3.10}$$

Thus η is a coboundary precisely when (3.10) holds, and so the cohomology of U is

$$H^1 = C^\infty(I)^3 / \{a_{EF}(m) + a_{FG}(m) + a_{GE}(m) = 0\}. \tag{3.11}$$

Lemma 3.3. The quotient (3.11) is isomorphic to \mathbb{C} .

Proof. Map $C^\infty(I)^3 \rightarrow \mathbb{C}$ by $\phi(f, g, h) = f(m) + g(m) + h(m)$, where m is the (unique) integer in I . This surjective homomorphism gives the desired isomorphism. \square

Finally, note that the condition $e^{2\pi it} = 1$ occurs precisely at the Bohr-Sommerfeld leaves. Therefore, we have proved the first part of the following result. The second statement follows since for $k > 1$, there are no $(k+1)$ -fold intersections in this cover.

Proposition 3.4. Let U be a band around the cylinder. Then the sheaf cohomology of U with respect to the cover $\mathcal{E}_3 = \{E, F, G\}$ is trivial if U does not contain a Bohr-Sommerfeld leaf. If U contains one Bohr-Sommerfeld leaf, its cohomology is

$$H_{\mathcal{E}_3}^k(U; \mathcal{J}) \cong \begin{cases} \mathbb{C} & k = 1 \\ 0 & k \neq 1 \end{cases}$$

This is precisely what we expect from Śniatycki's results.

Lemma 3.5. Let \mathcal{E}_k be a cover of the band similar to \mathcal{E}_3 , but with k sets instead of 3. Then the cohomology calculated with respect to \mathcal{E}_k is the same as that calculated with respect to \mathcal{E}_3 .

Proof. In this case, the same argument as for the cover \mathcal{E}_3 applies; the only difference is that we have k equations instead of 3 in (3.3) and (3.7), and the matrix (3.9) is a $k \times k$ matrix. Its determinant is the same, however, (up to sign) and we obtain the same conclusion as in Proposition 3.4 for cohomology with respect to the cover \mathcal{E}_k . \square

3.3 Brick wall covers

Eventually (see 3.5 below), we will find the cohomology of a band U by breaking it up into “sub-bands,” finding the cohomology of each one by hand, and then piecing them together using Mayer-Vietoris. If each band has a cover of the form \mathcal{E}_k , the cover induced on their intersection will be similar to \mathcal{E}_k but with two “layers,” and thus the calculation of section 3.2 is not sufficient. In this section, we define a type of covering we call a “brick wall,” which looks like \mathcal{E}_k but with more layers, and compute the cohomology of a band with this type of cover.

Later (in 3.5), we will take the direct limit over all covers of U (see 2.2) to find the actual sheaf cohomology. We will use Lemma 2.4 and show that the set of brick wall covers is cofinal in the set of all covers of U (see Lemma 3.14 below), and so it will be sufficient for all our purposes to consider only brick wall covers.

It will be enough to consider covers with only two layers of bricks, since the cohomology of covers with more layers of bricks can be found by piecing together two layers at a time, using Mayer-Vietoris.

Definition. A *brick wall* cover of a band in the cylinder (or, more generally, of any rectangle) is a finite covering by open $t\text{-}\theta$ rectangles (“bricks”), satisfying the following properties:

- The rectangles can be partitioned into sets (“layers”) so that all rectangles in one set cover the same interval of t values (“All bricks in the same layer have the same height”);
- Each brick contains points that are not in any other brick; and
- There are no worse than triple intersections, i.e., the intersection of two bricks in one layer does not meet the intersection of two bricks in either of the two adjoining layers.

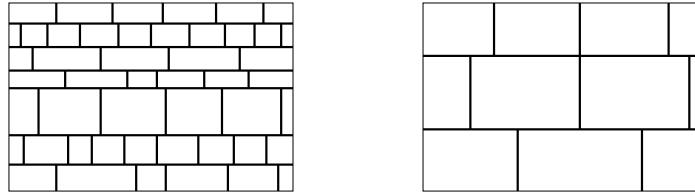


Figure 2: A brick wall cover, and one which is not

Note that we do not require that the number of bricks be the same in each layer, nor that the layers have the same height, nor that the bricks within one layer have the same width. See Figure 2, where we have “unrolled” the band, and where thick lines indicate intersections. Later we will allow brick walls with countably many layers, although the number of bricks in each layer will still be finite.

Let $U = I \times S^1$ be a band around the cylinder, which we cover by a brick wall of two layers. Let the top layer have m bricks A_1 through A_m , and the bottom layer have n bricks B_1 through B_n . (We choose our numbering of the A 's and B 's so that B_1 meets both A_1 and A_m .) Denote this covering by \mathcal{B}_n^m . Let I_A and I_B denote the intervals of t values which are covered by the A and B layers respectively. We also assume that U contains at most one Bohr-Sommerfeld leaf, and that this leaf is not contained in the intersection between the two layers. (It is possible to perform the calculation without this assumption. However, this case is sufficient for our purposes—since eventually we'll be taking finer and finer covers, we can always arrange that the Bohr-Sommerfeld leaves avoid the intersections between brick layers—and avoids some complications in the argument.) The most challenging part of this calculation is the bookkeeping, so we will set out our notational conventions at the outset.

A 0-cochain is given by a leafwise flat section on each brick in the cover, which as we showed in 3.1 is of the form $f(t)e^{it\theta}$ for some smooth function f .

We will denote the functions corresponding to the set A_j and B_k by a_j and b_k , respectively, so that for example $a_k(t)e^{it\theta}$ is the element of the cochain defined on A_k .

A 1-cochain is given by a section on each intersection of two bricks. The functions corresponding to the intersections $A_i \cap A_j$ and $B_k \cap B_l$ will be denoted by a_{ij} and b_{kl} respectively. The function corresponding to the set $A_j \cap B_k$ will be denoted by c_{jk} . Thus, the a_{ij} 's and b_{kl} 's give the parts of the 1-cochain defined on the intersections *within* one layer of the brick wall, while the c_{jk} 's are on the intersections *between* the layers. In particular, the a 's are functions of t defined on I_A , the b 's are defined on I_B , and the c 's are defined on $I_{AB} = I_A \cap I_B$.

As we noted before, the order in which sets are written in intersections is important in Čech cohomology—it doesn't matter how we do it, as long as we're consistent—and so we set out our conventions here. We will write sets with smaller index before those with larger index (wrapping around, so that n is considered “less” than 1), and sets on the A layers will be written before sets in the B layer. (Thus, we would write A_2A_3 , A_mA_1 , and $A_mA_1B_1$.) Also, on an intersection, we will use the θ coordinate from the set written *first* in the intersection by these conventions. Finally, the “branches” of θ will be chosen so that $\theta_{A_m} = \theta_{A_1} + 2\pi$ on A_mA_1 , $\theta_{B_n} = \theta_{B_1} + 2\pi$ on B_nB_1 , $\theta_{B_1} = \theta_{A_1}$ on A_1B_1 , and the θ coordinates on all other double intersections agree.

As in section 3.2, $H^0 = 0$ because there are no global sections, and $H^j = 0$ for $j \geq 3$ because this cover has at most triple intersections. We start with H^1 —as with the simpler cover, this is where all the interesting things happen—and deal with H^2 later.

Suppose we're given a 1-cocycle. This is a collection of $a_{ij}(t)$, $b_{kl}(t)$, and $c_{jk}(t)$, for all possible intersections of the A 's and B 's,⁸ satisfying certain conditions, which we'll deal with in a moment. We seek a_j , b_k defined on each A_j and B_k which make up a 0-cochain whose coboundary is our $\{a_{ij}, b_{kl}, c_{jk}\}$.

Just looking at the bricks within one layer, the situation is identical to the cover \mathcal{E}_k considered in the previous section, and we get a system of equations like (3.7). For the A layer, we get m equations in the m unknown functions a_1, \dots, a_m defined on I_A , which, after applying the relationships between

⁸ Note that the a_{ij} and b_{kl} will actually be

$$a_{i(i+1)} \quad \text{and} \quad b_{k(k+1)}, \quad 1 \leq i \leq m-1, \quad 1 \leq k \leq n-1,$$

plus a_{m1} and b_{n1} . The c_{jk} , on the other hand, will be defined for all pairs (j, k) for which A_j and B_k intersect; it is not possible a priori to say which pairs exist, but as discussed in a moment, there will be $m+n$ of the c_{jk} .

the various θ coordinates and cancelling common factors of $e^{it\theta}$, give the m equations

$$\begin{aligned} a_{12} &= a_2 - a_1 \\ a_{23} &= a_3 - a_2 \\ &\vdots \\ a_{m1} &= a_m - e^{2\pi it} a_1 \end{aligned} \tag{3.12}$$

as functions of t defined on I_A . Provided these equations are consistent, they uniquely determine a_1, \dots, a_m from the a_{ij} (by the same linear algebra argument as for the cover \mathcal{E}_k in Lemma 3.5). As in that case, these equations will be consistent provided $e^{2\pi it} \neq 1$ on I_A .

A similar set of n equations:

$$\begin{aligned} b_{12} &= b_2 - b_1 \\ &\vdots \\ b_{n1} &= b_n - e^{2\pi it} b_1 \end{aligned} \tag{3.13}$$

holds for the n functions b_k on I_B , which, again provided $e^{2\pi it} \neq 1$ (on I_B), uniquely determine the b_k from the b_{kl} .

Thus, all of the functions making up our 0-cochain are already determined just from the elements of the 1-cocycle which only exist within one layer. However, we also have a number of equations coming from the intersections between the layers, which need to be taken into account.

First of all, note that there are $n + m$ double intersections between the layers. This can be seen easily from Figure 3, which shows the view of the middle of a typical two-layer brick wall cover. Start at one point and go around the cylinder, counting double intersections. A new one will be added to the count every time we pass one of the vertical lines, i.e. the intersection between two bricks in the same layer. Since there are m bricks in one layer and n in the other, there are a total of $m + n$ vertical lines, and thus $m + n$ double intersections between bricks in different layers. (A similar argument shows that there are $m + n$ triple intersections in this cover.)

On each such double intersection, there is a c_{jk} which must satisfy

$$c_{jk} = b_k - a_j \tag{3.14}$$

as functions of t on I_{AB} , and so we have $m + n$ c_{jk} 's and $m + n$ such equations.



Figure 3: The intersection of two brick layers

We also have relations among the a_{ij} , b_{ij} , and c_{jk} , coming from the fact that they make up a cocycle. These equations come from the triple intersections of sets in the cover. Since there are $m + n$ triple intersections, we have $m + n$ equations, which are of the form

$$\begin{aligned} c_{(i+1)k} - c_{ik} + a_{i(i+1)} &= 0 && \text{from } A_i A_{i+1} B_k \\ &\quad \text{or} \\ b_{k(k+1)} - c_{ik} + c_{i(k+1)} &= 0 && \text{from } A_i B_k B_{k+1} \end{aligned} \tag{3.15}$$

depending on whether the intersection has two A 's or two B 's. (This is ignoring, for the moment, when there are factors of $e^{2\pi it}$ to worry about, which only happens near A_1 and B_1 .) These are the cocycle conditions, which are again equalities of functions of t defined on I_{AB} .

Essentially, we have $m + n$ extra equations that the a_i and b_j must satisfy, but we have $m + n$ conditions that the extra equations satisfy, which is enough to cancel each other out. In what follows we show more detail, but this is the essential idea.

Re-write equations (3.15) as

$$c_{(i+1)k} - c_{ik} = -a_{i(i+1)} \tag{3.16a}$$

$$c_{ik} - c_{i(k+1)} = b_{k(k+1)} \tag{3.16b}$$

Think of the a_{ij} and b_{kl} as being given, and these $m + n$ equations as defining the c_{jk} in terms of them. Since there are $m + n$ unknowns (the c_{jk}), provided they are consistent and not underdetermined, they define the c_{jk} uniquely. This will be the case provided the determinant of the matrix of coefficients of the system is not zero.

Order the c_{jk} by just going around the central part of the band. From Figure 3, we can see that each double intersection meets two triple intersections; thus, each c_{jk} appears in two (successive) equations in (3.16): once with its predecessor, and once with its successor. If we write all of the equations in the form given in (3.16), then a given c_{jk} has a positive sign when it appears with its predecessor, and a negative sign when it appears with its successor, and these are its only appearances. Thus the coefficient matrix for the system

(3.16) will have mostly zeros, except each row will have one -1 and one $+1$ in adjacent entries.

Finally, there are two places where we need to take into account the change in θ coordinates, which happens between A_m and A_1 and between B_n and B_1 . The two corresponding cocycle equations are

$$c_{11}e^{it\theta_{A_1}} - c_{m1}e^{it\theta_{A_m}} + a_{m1}e^{it\theta_{A_m}} = 0 \quad \text{on } A_mA_1B_1 \quad (3.17a)$$

$$b_{n1}e^{it\theta_{B_n}} - c_{k1}e^{it\theta_{A_k}} + c_{kn}e^{it\theta_{A_k}} = 0 \quad \text{on } A_kB_nB_1 \quad (3.17b)$$

(where A_k is the brick in the A layer that straddles the intersection of B_1 and B_n).

Using the fact that $\theta_{A_1} = \theta_{A_m} - 2\pi$, we rewrite (3.17a) as

$$c_{11}e^{it\theta_{A_m}-2\pi it} - c_{m1}e^{it\theta_{A_m}} + a_{m1}e^{it\theta_{A_m}} = 0 \quad \text{on } A_mA_1B_1$$

which yields

$$c_{11}e^{-2\pi it} - c_{m1} = -a_{m1}. \quad (3.18)$$

As for (3.17b), since $\theta_{B_n} = \theta_{A_k}$, we can cancel the $e^{it\theta}$ terms immediately to get

$$c_{kn} - c_{k1} = b_{n1}. \quad (3.19)$$

Therefore, the coefficient matrix for the system (3.16) can be put in the form:

$$\begin{bmatrix} -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & & & \ddots & & \\ 0 & 0 & 0 & \cdots & -1 & 1 \\ e^{-2\pi it} & 0 & 0 & \cdots & 0 & -1 \end{bmatrix} \quad (3.20)$$

By expanding along the bottom row, we can see that this matrix has determinant $e^{-2\pi it} - 1$, and thus the equations (3.15) have a unique solution for the c_{jk} provided, as usual, that $e^{2\pi it} \neq 1$ on I_{AB} —which is precisely what we are assuming.

Finally, it is straightforward to check that $c_{jk} = b_k - a_j$, where the a_j and b_k are the ones found already, gives a solution to (3.15).

The upshot of all of this is that the a_i and b_k are determined entirely by the parts of the cocycle defined on the intersections within one layer, namely the a_{ij} and the b_{kl} . The parts of the cocycle defined on the intersections between the layers (the c_{jk}) don't have any effect on the a_i and b_k , because of the cocycle conditions. Thus, given a 1-cocycle, provided that $e^{2\pi it}$ is never

1 on the band, it is the coboundary of a 0-cochain, and so the first Čech cohomology is zero.

Now consider the case when $e^{2\pi it} = 1$ somewhere on the band. If $e^{2\pi it_0} = 1$ for some $t_0 \in I_A$, then by an elementary linear algebra argument, the equations (3.12) are only consistent if

$$a_{12}(t_0) + a_{23}(t_0) + \cdots + a_{m1}(t_0) = 0.$$

In this case, there is a unique solution for $\{b_1, \dots, b_n\}$ on I_B , but there is only a solution for $\{a_1, \dots, a_m\}$ if $a_1(t_0) + \cdots + a_m(t_0) = 0$. Since $m \notin I_{AB}$, the system (3.16) still has a unique solution, which is compatible with the solutions for the a 's and b 's by the cocycle conditions. Thus, in this case, H^1 will be given by

$$H^1 = Z^1 / \{a_1(t_0) + \cdots + a_m(t_0) = 0\} \quad (3.21)$$

where Z^1 is the space of 1-cocycles.

Lemma 3.6. The set of 1-cocycles Z^1 is isomorphic to

$$C^\infty(I_A)^m \oplus C^\infty(I_B)^n,$$

provided there is no integer in I_{AB} .

Proof. As discussed above, if we think of the $\{a_{ij}\}$ and $\{b_{kl}\}$ as being given, we can view the equations (3.15) as defining the functions $\{c_{jk}\}$ in terms of the a 's and b 's; as noted above, this system will have a unique solution if $e^{2\pi it} \neq 1$ on I_{AB} . Thus, specifying a 1-cocycle amounts to giving the m functions $\{a_{12}, a_{23}, \dots, a_{m1}\}$ on I_A and the n functions $\{b_{12}, b_{23}, \dots, b_{n1}\}$ on I_B . Thus $Z^1 \cong C^\infty(I_A)^m \oplus C^\infty(I_B)^n$. \square

Lemma 3.7. The quotient (3.21) is isomorphic to \mathbb{C} .

Proof. Map $Z^1 \cong C^\infty(I_A)^m \oplus C^\infty(I_B)^n \rightarrow \mathbb{C}$ via

$$(f_1, \dots, f_{m+n}) \mapsto f_1(t_0) + f_2(t_0) + \cdots + f_m(t_0).$$

As in the argument in Lemma 3.3 (Section 3.2), this is a surjective homomorphism, and gives the desired quotient. \square

On the other hand, if $e^{2\pi it} = 1$ on I_B , the same argument applies to the b_{kl} and gives the same result for the cohomology.

To sum up, we have shown the $k = 0$ and $k = 1$ cases of the following:

Lemma 3.8. Let U be a band around the cylinder containing at most one Bohr-Sommerfeld leaf, with a brick wall cover \mathcal{B}_n^m of two layers. Assume that the Bohr-Sommerfeld leaf is contained in at most one layer of bricks. The sheaf cohomology of U with respect to the cover \mathcal{B}_n^m is 0 if U contains no Bohr-Sommerfeld leaf, and

$$H_{\mathcal{B}_n^m}^k(U; \mathcal{J}) \cong \begin{cases} \mathbb{C} & k = 1 \\ 0 & k \neq 1 \end{cases} \quad (3.22)$$

if it contains one Bohr-Sommerfeld leaf.

Proof of $H^{\geq 2}$. Unlike the simpler cover of section 3.2, we do have triple intersections in this cover, and so it is not immediate that $H^2 = 0$. However, it will not take long to dispose of this calculation.

A 2-cochain is one section for each triple intersection, which again can be represented by a smooth function of t on the appropriate interval. Since all the triple intersections lie along the intersection of the two layers, these functions will be defined on I_{AB} . Denote these functions by f 's, and call them $f_{j(j+1)k}$ for the intersections $A_j A_{j+1} B_k$ and $f_{jk(k+1)}$ for $A_j B_k B_{k+1}$.

Since there are no quadruple intersections, there are no 3-cochains, and so every 2-cochain is a cocycle. Thus, to show that H^2 is trivial, we need to show that every 2-cochain is a coboundary. So suppose we are given the f 's making up a 2-cochain.

There are $m + n$ triple intersections, as noted above, and thus $m + n$ f 's. On each intersection we have an equation of the form

$$c_{(j+1)k} - c_{jk} + a_{j(j+1)} = f_{j(j+1)k} \quad (3.23)$$

or

$$b_{k(k+1)} - c_{jk} + c_{j(k+1)} = f_{jk(k+1)} \quad (3.24)$$

This gives $m + n$ equations in the $2(m + n)$ unknowns (the a 's, b 's, and c 's) and so has infinitely many solutions. (The equations are clearly linearly independent and so consistent.) Thus every 2-cocycle is a coboundary, and $H^2 = 0$.

Finally, $H^k = 0$ for $k \geq 3$ since there are no $k + 1$ -fold intersections. \square

3.4 Mayer-Vietoris

In this section we apply a Mayer-Vietoris type argument to find the cohomology of the union of two open sets. The argument follows that given in [BT], Section I.2.

Let $U = S^1 \times I_U$ and $V = S^1 \times I_V$ be bands around M , where I_U and I_V are intervals. Fix a covering $\mathcal{A} = \{A_\alpha\}$ of M . This induces coverings on U , V , and $U \cap V$ (which we will still denote by \mathcal{A}). In what follows, we compute sheaf cohomology always with respect to these covers.

Consider the sequence of Čech cochain complexes

$$0 \rightarrow C_{\mathcal{A}}^*(U \cup V, \mathcal{J}) \xrightarrow{q} C_{\mathcal{A}}^*(U, \mathcal{J}) \oplus C_{\mathcal{A}}^*(V, \mathcal{J}) \xrightarrow{r} C_{\mathcal{A}}^*(U \cap V, \mathcal{J}) \rightarrow 0 \quad (3.25)$$

For ease of notation, use α for $\alpha_0 \cdots \alpha_k$, so that $A_\alpha = A_{\alpha_0 \cdots \alpha_k}$. The map q restricts a component of a cochain f_α on A_α to $U \cap A_\alpha$ and $V \cap A_\alpha$, respectively. The map r is defined as follows:

$$(r(f, g))_\alpha = f_\alpha - g_\alpha \quad \text{on } U \cap V \cap A_\alpha.$$

Proposition 3.9. The sequence (3.25) is exact.

Proof. By construction $rq = 0$, so we wish to show $\ker r \subset \text{im } q$. Suppose $(f, g) \xrightarrow{r} 0$. This means that for each α , $f_\alpha = g_\alpha$ on $A_\alpha \cap U \cap V$. By the first axiom of sheaves (see 2.2), there exists a function h_α on $A_\alpha \cap (U \cup V)$ which restricts to f_α and g_α on the appropriate sets. The collection of h_α for each α gives a cochain h with $q(h) = (h|_U, h|_V) = (f, g)$. This shows exactness at the middle.

Exactness at the left end merely requires that q be injective, namely that if $f|_{U \cap A_\alpha} = g|_{U \cap A_\alpha}$ and $f|_{V \cap A_\alpha} = g|_{V \cap A_\alpha}$, then $f = g$ on $(U \cup V) \cap A_\alpha$. This is the second axiom of sheaves.

Exactness at the right end requires that r be surjective. This is the most interesting part, as we don't have any sheaf axioms left; the argument in [BT] uses partitions of unity, which do not exist for all sheaves. In our case, we have partitions of unity in the t directions, which is sufficient.

Explicitly, let A be an element from the cover \mathcal{A} . An element of \mathcal{J} looks like $h = a(t)e^{it\theta}$; suppose such an element is given on $A \cap U \cap V$. Let $\rho_U(t), \rho_V(t)$ be two functions on \mathbb{R} supported on I_U and I_V respectively, with $\rho_U + \rho_V \equiv 1$. Then $\rho_V h$ is a section over U , and $-\rho_U h$ is a section over V , both of which are in \mathcal{J} , since they are of the form (smooth function) $\times e^{it\theta}$. Furthermore, $\rho_V h - (-\rho_U h) = h$ on $U \cap V \cap A$. Thus $\rho_V h \oplus (-\rho_U h)$ maps to h under r .

Therefore r is surjective, and the sequence (3.25) is exact. \square

Proposition 3.10 (Band Sheaf Mayer-Vietoris). Let U and V be bands in M . There is a long exact sequence of sheaf cohomology

$$\begin{aligned} \cdots &\rightarrow H^k(U \cup V, \mathcal{J}) \rightarrow H^k(U, \mathcal{J}) \oplus H^k(V, \mathcal{J}) \rightarrow \\ &\rightarrow H^k(U \cap V, \mathcal{J}) \rightarrow H^{k+1}(U \cup V, \mathcal{J}) \rightarrow \cdots \end{aligned} \quad (3.26)$$

This sequence holds both for actual sheaf cohomology, and also for cohomology computed with respect to a particular cover.

Proof. For each covering \mathcal{A} of M , there is a short exact sequence of cochain complexes (3.25), computed with respect to \mathcal{A} . Each induces a long exact sequence in cohomology, as usual, which is (3.26) with respect to the cover \mathcal{A} , and so the cover-specific result is shown.

Next, the sheaf cohomology of M is the direct limit, over the set of open covers of M , of the cohomology with respect to each cover. The maps in the directed system of open covers are just the restriction of sheaf elements to smaller sets in a finer cover (see 2.2). Since the maps in the directed system of open covers are just restrictions, which commute with the coboundary maps in the cochain complexes, they induce transformations of cochain complexes. By Lemma 2.6, this induces a map between the corresponding long exact sequences of cohomology, namely, the Mayer-Vietoris sequences of cohomology with respect to the particular covers. Thus we have, in essence, a directed system of long exact sequences of the form (3.26), *with respect to the particular covers*.

The maps in this directed system, which are restrictions, commute with the maps in the long exact sequences, which are also restrictions or subtractions (defined following (3.25)). Hence, by Lemma 2.5, the exactness passes to the direct limit, and thus we have the sequence (3.26) in actual cohomology. \square

Corollary 3.11. Let U and V be bands in M , such that $U \cap V$ does not contain a Bohr-Sommerfeld leaf. Let \mathcal{B} be a brick wall cover of $U \cup V$, which restricts to a brick wall cover on $U \cap V$ with only two layers. Then

$$H_{\mathcal{B}}^k(U \cup V, \mathcal{J}) \cong H_{\mathcal{B}}^k(U, \mathcal{J}) \oplus H_{\mathcal{B}}^k(V, \mathcal{J}). \quad (3.27)$$

(The case we're most concerned with is $k = 1$. It is true for other values of k as well, though in this case both sides are 0.)

Proof. We have the sequence (3.26) with respect to the cover \mathcal{B} . Now $H_{\mathcal{B}}^{k-1}(U \cap V, \mathcal{J})$ and $H_{\mathcal{B}}^k(U \cap V, \mathcal{J})$ are both zero for all values of k , by the results in Section 3.3: if $k \neq 1$ then $H_{\mathcal{B}}^k(U \cap V, \mathcal{J}) = 0$ automatically, and if $k = 1$ then it follows from the assumption that $U \cap V$ contains no Bohr-Sommerfeld leaf. Therefore, the map in (3.26)

$$H_{\mathcal{B}}^k(U \cup V, \mathcal{J}) \rightarrow H_{\mathcal{B}}^k(U, \mathcal{J}) \oplus H_{\mathcal{B}}^k(V, \mathcal{J})$$

is an isomorphism. \square

3.5 Refinements and covers: scaling the brick wall

In this section we show that the cohomology computed in the preceding sections is the actual sheaf cohomology of U .

Theorem 3.12. Let U be a band in the cylinder M . Then

$$H^k(U; \mathcal{J}) \cong \begin{cases} \mathbb{C}^\nu & k = 1 \\ 0 & k \neq 1 \end{cases} \quad (3.28)$$

where ν is the number of Bohr-Sommerfeld leaves contained in U , and where the cohomology is the actual sheaf cohomology.

The proof of this theorem uses a couple of technical lemmas.

Lemma 3.13. The cohomology of a band which contains no Bohr-Sommerfeld leaf, with respect to a brick wall cover, is trivial, even if the brick wall has countably many layers.

Proof. Suppose a brick wall covering \mathcal{B} is given.

We first show that $H^1 = 0$. This requires that, given a 1-cocycle η , we can find a 0-cochain β whose coboundary is η .

Index the layers by some appropriate set of integers, and denote the k^{th} layer by R^k , made up of m_k bricks. Then η is a collection of m_k functions defined on the intersections between bricks in each layer R^k , plus a number of functions defined on the intersections between adjacent layers. The 0-cochain β , on the other hand, is a collection of m_k functions $\beta_1^k, \dots, \beta_{m_k}^k$ on each of the bricks $B_1^k, \dots, B_{m_k}^k$ in R^k , for all k .

As we saw in section 3.3, however, the functions $\beta_1^k, \dots, \beta_{m_k}^k$ are uniquely determined from the intersections within the layer B_k . We also saw that the cocycle conditions guarantee that these solutions are consistent with the requirements that come from the parts of η defined on the intersections between bricks in different layers. More briefly, the β_j^k 's are defined entirely by the parts of η living on the k^{th} layer, and the layers above and below don't interfere.

Thus, given η , β is uniquely determined on each brick, and thus uniquely determined as a cochain. Thus η is a coboundary, and H^1 is zero.

The argument for the cohomology in other dimensions is similar. \square

The following is a standard result (for example, [Dg], Theorem XI.4.5):

Lemma (Lebesgue's Number Lemma). Given a covering of a compact metric space X , there is a number δ such that any subspace of X of diameter less than δ is contained in one of the sets of the cover.

Lemma 3.14. Any open cover \mathcal{A} of a band U has a refinement which is a brick wall cover (with possibly countably many layers of bricks). (Recall that one cover \mathcal{B} is a *refinement* of another cover \mathcal{A} if every element of \mathcal{B} is contained in some element of \mathcal{A} .)

Proof. Assume a cover \mathcal{A} of $U = I \times S^1$ is given.

Divide U up into a countable union of *closed* bands $\{R^k\}_{k=-\infty}^\infty$, by dividing I up into a union of closed intervals, with the properties

- each R^k overlaps its two neighbours R^{k-1} and R^{k+1} by some positive amount (i.e. their interiors overlap);
- R^k does not intersect any R^j other than its two neighbours; and
- the union of the R^k equals U .

The covering \mathcal{A} induces a covering of each layer R^k .

Starting with R^0 and proceeding inductively in both directions, choose a covering $\{B_1^k, \dots, B_{n_k}^k\}$ of R^k , with each B_j^k a *closed* rectangle, so that

- each B_j^k is contained in some element of \mathcal{A} (which, if the B_j^k are chosen small enough, is possible by Lebesgue's number lemma—if necessary, divide R^k into more layers),
- the collection $\mathcal{B}_k = \{\overset{\circ}{B}_1^k, \dots, \overset{\circ}{B}_{n_k}^k\}$ form a layer of bricks (in terms of their overlaps), where \circ denotes interior, and
- the coverings $\mathcal{B}_{-k}, \dots, \mathcal{B}_k$ form a brick wall (mostly, this amounts to choosing the overlaps between bricks in one layer to avoid the overlaps between bricks in the neighbouring layers).

Then the cover $\mathcal{B} := \bigcup_{k=-\infty}^\infty \mathcal{B}_k$ is a brick wall covering of U which is a refinement of \mathcal{A} . (The fact that the partial cover is a brick wall at each stage ensures that the entire cover is a brick wall; the condition that the interiors of the R^k 's overlap ensures that there are no “gaps” in the cover, and the fact that the R^k 's cover U guarantees that the $\overset{\circ}{B}_j^k$'s cover all of U .) \square

Proof of Theorem 3.12. Let \mathcal{B} be a brick wall covering of U , with possibly countably many layers, such that no Bohr-Sommerfeld leaf lies in more than one layer of bricks.

Denote by B_n the layer containing the n^{th} Bohr-Sommerfeld leaf (i.e. the one corresponding to the integer n), and let V_n be the union of all

layers of bricks between (but not including) B_n and B_{n+1} . Note that B_n and V_n are bands, and V_n contains no Bohr-Sommerfeld leaf. Then by Lemma 3.13, $H_B^*(V_n; \mathcal{J}) = 0$, while from the calculations in sections 3.2 and 3.3, $H_B^*(B_n; \mathcal{J}) \cong \mathbb{C}$ for each n appearing.

Since U is a finite union of B_n 's and V_n 's, $H_B^*(U; \mathcal{J})$ is the finite sum of the cohomologies of B_n and V_n , by the cover-specific Mayer-Vietoris. Thus we have

$$H_B^k(U; \mathcal{J}) \cong \begin{cases} \mathbb{C}^\nu & k = 1 \\ 0 & k \neq 1 \end{cases} \quad (3.29)$$

for any brick wall covering \mathcal{B} .

However, by Lemma 3.14, any covering of U has a refinement which is a brick wall. In the language of direct limits, the set of brick wall coverings is cofinal in the directed system of coverings used to calculate sheaf cohomology. (See Lemma 2.4.) Since the cohomology computed using any brick wall cover is the same, cofinality means that the actual sheaf cohomology is the one computed with these covers, and thus we have (3.28). \square

Corollary 3.15. Let U and V be bands in M , such that $U \cap V$ does not contain a Bohr-Sommerfeld leaf. Then

$$H^k(U \cup V, \mathcal{J}) \cong H^k(U, \mathcal{J}) \oplus H^k(V, \mathcal{J}). \quad (3.30)$$

in actual sheaf cohomology.

Proof. This is just Corollary 3.11, without the dependence on the cover. Since we now know that $H^{k-1}(U \cap V, \mathcal{J})$ and $H^k(U \cap V, \mathcal{J})$ are both zero for all values of k , without the cover dependence, the same argument as in Corollary 3.11 shows that

$$H^k(U \cup V, \mathcal{J}) \cong H^k(U, \mathcal{J}) \oplus H^k(V, \mathcal{J}).$$

\square

4 The complex plane

The second model space we consider is the complex plane \mathbb{C} . In this section we describe the set-up of the model space, and compute the the sheaf cohomology of \mathbb{C} by hand. The result we obtain is a little surprising, in that the count over Bohr-Sommerfeld leaves excludes the origin. The heart of the surprising result is Propositions 4.3 and 4.6.

The coordinates we use on \mathbb{C} are (s, ϕ) , where (r, ϕ) are standard polar coordinates and $s = \frac{1}{2}r^2$. In these coordinates, the standard symplectic form

has the expression $\omega = ds \wedge d\phi$. (Note that $d\phi$ is not defined at $s = 0$, but this form extends smoothly there.)

The plane is equipped with a singular real polarization given by the distribution $P = \text{span}\{\frac{\partial}{\partial\phi}\}$, which is integrable. Its integral manifolds are the circles of constant s . Note that this is not quite a foliation, because the manifold with $s = 0$ is a point, while the others are circles, but it is a singular foliation, as in 2.4.

4.1 The sheaf of sections flat along the leaves

Let \mathbb{L} be the trivial line bundle $\mathbb{C} \times \mathbb{C}$. The symplectic form ω is exact ($\omega = d(s d\phi)$), and so the connection defined in the canonical trivialization of \mathbb{L} by

$$\nabla_X \sigma = d\sigma(X) - \sigma is d\phi(X)$$

(where $\sigma: \mathbb{C} \rightarrow \mathbb{C}$) is a prequantization connection.

As before, we denote by \mathcal{J} the sheaf of sections flat along the leaves.

Proposition 4.1. The sections which are flat along the leaves are of the form $a(s)e^{is\phi}$, for arbitrary smooth functions a .

Proof. The argument is very similar to the argument in section 3.1, with t replaced by s . A section of \mathbb{L} over U is given by a map $\sigma: U \rightarrow \mathbb{C}$. It will be flat along the leaves if

$$0 = \nabla_X \sigma = X(\sigma) - \sigma is d\phi(X) \quad \forall X \in P.$$

Since P is the span of $\frac{\partial}{\partial\phi}$, this is the same as

$$0 = \nabla_{\frac{\partial}{\partial\phi}} \sigma = \frac{\partial\sigma}{\partial\phi} - \sigma is,$$

i.e.

$$\frac{\partial\sigma}{\partial\phi} = is\sigma.$$

Thus, the sections flat along the leaves are those of the form

$$\sigma = a(s)e^{is\phi} \tag{4.1}$$

for a some smooth function of s . □

This calculation applies anywhere but at the origin ($s = 0$), as $\frac{\partial}{\partial\theta}$ is not defined there. However, as we will see (Prop 4.3), this is enough to determine the value of σ at 0.

Lemma 4.2. The Bohr-Sommerfeld leaves on \mathbb{C} are precisely the circles $\{s = k\}_{k \in \mathbb{N}}$. In other words, if s is not an integer, there is no (nonzero) flat section defined over all of ℓ_s .

Proof. This is exactly the same argument as in Section 3: in order for $e^{is\phi}$ to be defined on an entire leaf, that is, the entire range of ϕ from 0 to 2π , it is necessary for s to be an integer. \square

Proposition 4.3. If U is a small open disc centered at 0, then $\mathcal{J}(U) = 0$.

Proof. By Lemma 4.2, a leaf only possesses a global flat section if its s value is an integer. Since U is a disc, it is made up of the union of leaves. On every leaf which has a non-integer s value, and in particular, all those near the origin, a flat section must be zero. Since a section is continuous, it must be zero everywhere. Thus there are no flat sections over U other than the zero section. \square

Another way of stating this is: Any section must be 0 on any open disc around the origin on which it is defined. This implies that the stalk of \mathcal{J} over 0 is zero, but it is actually a stronger condition.

4.2 Cohomology

In this section we calculate $H^k(U, \mathcal{J})$ for certain open sets $U \subset \mathbb{C}$ by a similar procedure as in section 3.

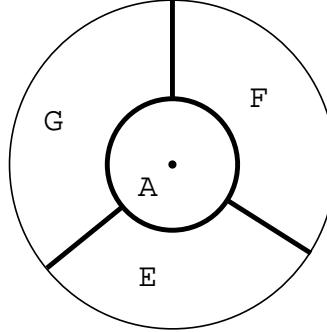
First of all, suppose U is an annulus centered at the origin, $\{(s, \phi) \mid r_0 < s < r_1, r_0 > 0\}$. This set is analogous to the band considered in section 3, and we have the same result for it:

Proposition 4.4. If U is an annulus centered at the origin which contains at most one Bohr-Sommerfeld leaf, its sheaf cohomology is

$$H^1(U, \mathcal{J}) = \begin{cases} \mathbb{C} & \text{if } U \text{ contains a Bohr-Sommerfeld leaf} \\ 0 & \text{if not} \end{cases} \quad (4.2)$$

and $H^k(U, \mathcal{J}) = 0$ for $k \neq 1$.

Proof. This is the same proof as for Prop 3.4. The set U and the elements of the sheaf have the same form here as there, with t replaced by s , and the same argument goes through word for word. \square

Figure 4: The cover \mathcal{A}

Next, consider the open set A which is a disc centred at the origin, surrounded by three sets E, F, G in a ring, none of which intersects a Bohr-Sommerfeld orbit. (See Figure 4.) Let \mathcal{A} denote the cover $\{A, E, F, G\}$, and $U = A \cup E \cup F \cup G$. We will calculate $H_{\mathcal{A}}^k(U; \mathcal{J})$, beginning with H^1 .

As before, we set out our bookkeeping conventions at the outset. We will use the same conventions as in 3.2, namely, that intersections will be written in the order EF, FG, GE . For intersections involving A , the A will always come first: AEF , etc. On an intersection, we will always use the ϕ coordinate from the set written *first* in these conventions, *unless that set is A*, in which case we use the second set (since there is no ϕ coordinate defined on all of A). We choose the ϕ coordinates so that $\phi_G = \phi_E + 2\pi$ on GE while the ϕ coordinates agree on the other intersections. Finally, let I_A denote the range of s -intervals covered by A , I_O (for ‘outer’) denote the range of s -intervals covered by E, F , and G , and $I_{AO} = I_A \cap I_O$.

Let $\alpha = \{\alpha_{VW}\}_{V,W=A,E,F,G}$ be a 1-cochain. The coboundary of α is given by

$$(\delta\alpha)_{VWX} = \alpha_{WX} - \alpha_{VX} + \alpha_{VW}.$$

Thus if α is a cocycle, $\delta\alpha = 0$ and we have the following equations:

$$\begin{aligned} h_{EFE}e^{is\phi_E} - h_{AFE}e^{is\phi_F} + h_{AEE}e^{is\phi_E} &= 0 && \text{on } AEF \\ h_{FG}e^{is\phi_F} - h_{AGE}e^{is\phi_G} + h_{AFE}e^{is\phi_F} &= 0 && \text{on } AFG \\ h_{GE}e^{is\phi_G} - h_{AEE}e^{is\phi_E} + h_{AGE}e^{is\phi_G} &= 0 && \text{on } AGE \end{aligned} \tag{4.3}$$

where the h 's are smooth functions of s defined on the appropriate interval.

Cancelling the factors of $e^{is\phi}$, we get

$$h_{EF} - h_{AF} + h_{AE} = 0 \quad (4.4a)$$

$$h_{FG} - h_{AG} + h_{AF} = 0 \quad (4.4b)$$

$$h_{GE} - e^{-2\pi is} h_{AE} + h_{AG} = 0 \quad (4.4c)$$

as functions of s , on I_{AO} .

Next, to see if α is a coboundary, we ask if there exists a 1-cochain β such that $\delta\beta = \alpha$. Unravelling the definitions, given $h_{VW}(s)$ satisfying equations (4.4), we seek elements of $\mathcal{J} \{g_V e^{is\phi_V}\}_{V=A,E,F,G}$, each defined on the corresponding set, such that

$$g_W e^{is\phi_W} - g_V e^{is\phi_V} = h_{VW} e^{is\phi} \quad \text{on } V \cap W \quad (4.5)$$

for all choices of V, W .

First, on AE , (4.5) would give

$$g_E e^{is\phi_E} - g_A e^{is\phi_A} = h_{AE} e^{is\phi_E} \quad \text{on } AE, \quad (4.6)$$

but by Prop. 4.3, there are no nonzero sections on A . So the term involving g_A is zero, and we get $g_E = h_{AE}$ on I_{AO} . Similarly, $g_F = h_{AF}$ and $g_G = h_{AG}$ (on I_{AO}). This determines β on all intersections with A .

For the definition of β outside of A , we need to define the g s on the rest of I_O .

Lemma 4.5. The functions g_E , g_F , and g_G extend to functions on all of I_O .

Proof. The functions h_{EF} , h_{FG} , and h_{GE} are defined on I_O ; the functions h_{AE} , h_{AF} , and h_{AG} are defined on I_{AO} . At present, we have defined g_E , g_F , and g_G only on I_{AO} .

Adding up the three equations (4.4a), (4.4b), and (4.4c) and rearranging, we obtain

$$(1 - e^{-2\pi is})h_{AE}(s) = -(h_{EF}(s) + h_{FG}(s) + h_{GE}(s)) \quad (4.7)$$

which is true on I_{AO} . But the functions on the right side are defined on all of I_O . Provided $e^{-2\pi is} \neq 1$ (which is true by our assumption that none of these sets contain a Bohr-Sommerfeld leaf), this gives that $h_{AE}(s)$ on I_{AO} is equal to a function that is defined on all of I_O . We can use this to define g_E on the rest of I_O .

Similarly, by adding $e^{-2\pi is} \cdot (4.4a)$, $(4.4b)$, and $(4.4c)$, we get a similar result for g_F ; the case for g_G is similar. For reference, we collect here the equations defining the g 's:

$$(1 - e^{-2\pi is})g_E = -(h_{EF} + h_{FG} + h_{GE}) \quad (4.8a)$$

$$(1 - e^{-2\pi is})g_F = -(e^{-2\pi is}h_{EF} + h_{FG} + h_{GE}) \quad (4.8b)$$

$$(1 - e^{-2\pi is})g_G = -(e^{-2\pi is}h_{EF} + e^{-2\pi is}h_{FG} + h_{GE}) \quad (4.8c)$$

□

These extensions define β on the rest of the sets E , F , and G .

Finally, there are additional conditions that β must satisfy, which arise from the intersections around the outside of the ‘ring,’ namely EF , FG , and GE . If we apply (4.5) to these intersections, we get the following conditions, which are required for $\delta\beta = \alpha$:

$$\begin{aligned} h_{EF}e^{is\phi_E} &= g_F e^{is\phi_F} - g_E e^{is\phi_E} && \text{on } EF \\ h_{FG}e^{is\phi_F} &= g_G e^{is\phi_G} - g_F e^{is\phi_F} && \text{on } FG \\ h_{GE}e^{is\phi_G} &= g_E e^{is\phi_E} - g_G e^{is\phi_G} && \text{on } GE \end{aligned} \quad (4.9)$$

Using the convention that $\phi_G = \phi_E + 2\pi$ and cancelling factors of $e^{is\phi}$ as before, we obtain

$$\begin{aligned} h_{EF} &= g_F - g_E \\ h_{FG} &= g_G - g_F \\ h_{GE} &= e^{-2\pi is}g_E - g_G \end{aligned} \quad (4.10)$$

on I_{AO} . However, the g s given by (4.8), satisfy these equations with no further restrictions (still assuming $e^{2\pi is} \neq 1$). For example, using (4.8b) and (4.8a),

$$\begin{aligned} g_F - g_E &= -\frac{1}{(1-e^{-2\pi is})}(e^{-2\pi is}h_{EF} + h_{FG} + h_{GE}) + \frac{1}{1-e^{-2\pi is}}(h_{EF} + h_{GE} + h_{FG}) \\ &= \frac{1}{1-e^{-2\pi is}}((1 - e^{2\pi is})h_{EF}) \\ &= h_{EF} \end{aligned}$$

as required.

(In fact, the required conditions (4.10) are the same as the cocycle conditions (4.4), except that the latter only apply on the intersections with A , while the above calculation applies on the entire domain of β . This is reminiscent of what we saw in 3.3 (around equations (3.15) and (3.16)), where the

conditions for compatibility of two sets of solutions were precisely the cocycle conditions.)

Therefore, given any cocycle $\alpha \in C^1$, there exists a $\beta \in C^0$ such that $\delta\beta = \alpha$, and hence every 1-cocycle is a coboundary. Thus, we have proved the following:

Theorem 4.6. For $U \in \mathbb{C}$ an open disc centred at 0, and \mathcal{A} the cover given in Figure 4, $H_{\mathcal{A}}^1(U; \mathcal{J}) = 0$.

The cohomology in all other dimensions is trivial:

Proposition 4.7. $H_{\mathcal{A}}^k(U; \mathcal{J}) = 0$ for $k \neq 1$ as well.

Proof. First consider $k = 0$. As noted in section 3.2, $H^0(U; \mathcal{J})$ is the space of global sections of \mathcal{J} over U . Since there are no such sections, H^0 is zero.

Next, consider H^2 .

A 2-cochain α is a collection $\{a_{AEF}e^{is\phi_E}, a_{AFG}e^{is\phi_F}, a_{AGE}e^{is\phi_G}\}$, and is automatically a cocycle.

A 1-cochain β is a collection $\{b_{VWE}e^{is\phi_U}\}_{U,V=E,F,G,A}$. Now α will be $\delta\beta$ if

$$\begin{aligned} b_{EFE}e^{is\phi_E} - b_{AFE}e^{is\phi_F} + b_{AEE}e^{is\phi_E} &= a_{AEF}e^{is\phi_E} && \text{on } AEF \\ b_{FGF}e^{is\phi_F} - b_{AGF}e^{is\phi_G} + b_{AFG}e^{is\phi_F} &= a_{AFG}e^{is\phi_F} && \text{on } AFG \\ b_{GEG}e^{is\phi_G} - b_{AEG}e^{is\phi_E} + b_{AGE}e^{is\phi_G} &= a_{AGE}e^{is\phi_G} && \text{on } AGE \end{aligned} \quad (4.11)$$

Using the convention that $\phi_G = \phi_E + 2\pi$ on GE , and cancelling factors of $e^{is\phi}$, we obtain:

$$\begin{aligned} a_{AEF} &= b_{EF} - b_{AF} + b_{AE} \\ a_{AFG} &= b_{FG} - b_{AG} + b_{AF} \\ a_{AGE} &= b_{GE} - e^{-2\pi is}b_{AE} + b_{AG} \end{aligned} \quad (4.12)$$

on I_{AO} . This is a system of 3 equations in 6 unknowns, and thus has many solutions. (It is easy to check that it is consistent.) There is the question of extendability: for example, a_{AGE} is only defined for $s \in I_{AO}$, while b_{EF} needs to be defined for $s \in I_O$. However, if we set b_{EF} , b_{FG} , and b_{GE} to zero, the resulting system in b_{AE} , b_{AF} , and b_{AG} is the same system as (3.8), which has a solution provided that $e^{2\pi is} \neq 1$. Since there are no Bohr-Sommerfeld leaves in any of the sets we are considering, this condition holds, and so the system (4.12) always has a solution. Thus, any cocycle is a coboundary, and the second cohomology is trivial.

Finally, for $k \geq 3$, there are no $(k+1)$ -fold intersections, and so there is no cohomology. \square

Proposition 4.8. The cohomology calculated with respect to this particular cover is the actual sheaf cohomology.

Proof. This is essentially the same argument as in Section 3.5. The only modification that we need is in the form of the cover, as a brick wall doesn't immediately apply to the complex plane. The types of covers we use are as follows: a brick wall covering (a brick wall in s, ϕ coordinates) of the set $\{s \neq 0\}$, plus a disc centred at zero. It is clear that any cover has a refinement of this form: away from the origin the same argument applies as with the brick wall, and around the origin we need only take a small enough disc. \square

4.3 Mayer-Vietoris

Proposition 4.9. If each of $U, V \subset \mathbb{C}$ is either an open annulus or an open disc, centered at the origin, then the sequence

$$\cdots \rightarrow H^1(U \cup V, \mathcal{J}) \rightarrow H^1(U, \mathcal{J}) \oplus H^1(V, \mathcal{J}) \rightarrow H^1(U \cap V, \mathcal{J}) \rightarrow H^2(U \cup V, \mathcal{J}) \rightarrow \cdots \quad (4.13)$$

as in Prop. 3.10, is exact.

Proof. Away from 0, the exact same argument as in section 3.4 applies, with t replaced by s . Near 0, we need another argument.

Assume U is a disc centered at 0, and V is an annulus centered at 0 which overlaps with U and whose closure does not contain 0. Consider the sequence of Čech cochain complexes computed with respect to some cover \mathcal{B} (which here may be any cover):

$$0 \rightarrow C_{\mathcal{B}}^*(U \cup V, \mathcal{J}) \xrightarrow{q} C_{\mathcal{B}}^*(U, \mathcal{J}) \oplus C_{\mathcal{B}}^*(V, \mathcal{J}) \xrightarrow{r} C_{\mathcal{B}}^*(U \cap V, \mathcal{J}) \rightarrow 0 \quad (4.14)$$

which is (3.25) from section 3.4 (where the definition of the maps r and q is given). Exactness at the left and the centre only relies on properties of sheaves, as described in the proof of Prop 3.9, and thus still hold in this case. Exactness on the right requires r to be surjective, namely, given some $h \in \mathcal{J}(B \cap U \cap V)$, we need $f \in \mathcal{J}(U)$ and $g \in \mathcal{J}(V)$ whose difference on $B \cap U \cap V$ is h . As in the proof of Prop 3.9, we take a partition of unity $(\rho_U(s), \rho_V(s))$ over the s -intervals covered by U and V , and let $f = \rho_V h$ on $U \cap B$, $g = \rho_U h$ on $V \cap B$. The only possible concern is that, if B contains the origin, f must be zero on any disc centred at 0 contained in $B \cap U$. Since \overline{V} does not contain 0, this means there is a disc around zero outside \overline{V} ; and since ρ_V is zero outside \overline{V} , f must therefore be zero on this disc, and therefore on any disc in $B \cap U$. \square

By repeated applications of Mayer-Vietoris, we obtain:

Theorem 4.10. Let $U \subset \mathbb{C}$ be either an annulus or a disc, centred at the origin. Then

$$H^1(U, \mathcal{J}) \cong \mathbb{C}^m$$

where m is the number of Bohr-Sommerfeld leaves, *excluding* the origin, contained in U .

5 Example: S^2

The simplest example of a toric manifold, which is ubiquitous⁹ in textbooks on symplectic geometry, is S^2 with an action of S^1 by rotations about the z -axis. In this case the moment map is simply the height function. (See Figure 5.) This example serves as a good illustration of the results of this paper: despite its simplicity, it contains the essential idea of our method.

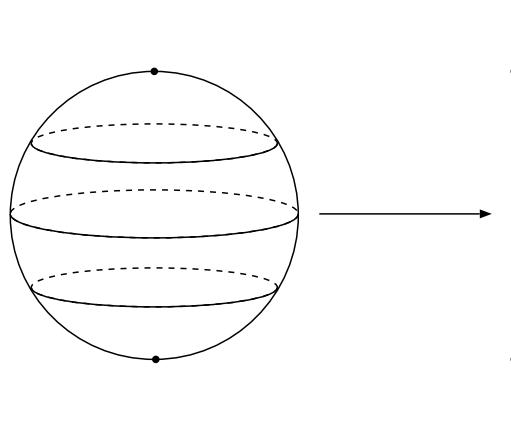


Figure 5: The moment map on the 2-sphere

In fact, once we draw the picture in Figure 6, further explanation is almost unnecessary. The orbits of the circle action are circles of constant height, plus the two singular orbits at the north and south poles. A neighbourhood of a circle orbit looks like a neighbourhood of a circle in the cylinder, and a neighborhood of one of the poles looks like a neighbourhood of the origin in

⁹ This is probably because it is the only toric manifold of dimension less than 4 (by Delzant's classification, see [CdS1]), and so it's the only one that can be drawn on a page.

\mathbb{R}^2 , as illustrated in Figure 6. We have determined the cohomology of each of these neighbourhoods in the preceding two sections, and we transfer these results over to the sphere.

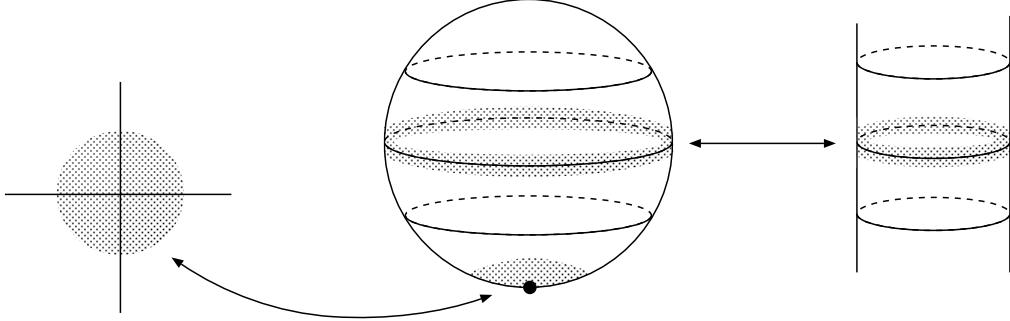


Figure 6: Neighbourhoods of orbits on the 2-sphere

More formally, let $M = S^2$, and ω_1 be the standard symplectic (area) form, normalized so that the total area of the sphere is 2π . Let \mathbb{L}_1 be the complex line bundle with Chern class 1, which is a prequantum line bundle for S^2 . Let $\mathbb{L}_k = \mathbb{L}_1^{\otimes k}$, which will then be a prequantum line bundle for S^2 with symplectic form $\omega_k := k\omega_1$, by the additivity of the Chern class.

As mentioned above, the height function μ is a moment map for the circle action, normalized so the total height of the sphere is 1; similarly, if we take the same circle action but use the symplectic form ω_k , then its moment map is k times the height function, which we denote by $\mu_k: S^2 \rightarrow \mathbb{R}$. The fibres give a singular Lagrangian fibration.

The Bohr-Sommerfeld leaves in this example are the leaves on which the moment map has an integer value, which we can see as follows. As in section 2.1.1, Bohr-Sommerfeld fibres are those with trivial holonomy, i.e. those leaves ℓ such that

$$\exp\left(i \int_{\ell} \Theta_k\right) = 1$$

where Θ_k is a potential 1-form for the connection on the prequantum line bundle \mathbb{L}_k . This will be true iff

$$\int_{\ell} \Theta_k = \int_{\Sigma} \omega_k \in 2\pi\mathbb{Z}$$

where Σ is a surface whose boundary is ℓ . Taking Σ to be the “bottom cap,” the set of points with height less than ℓ , then ℓ will be Bohr-Sommerfeld iff k

times the area of Σ is an integer multiple of 2π — that is, iff $\mu_k(\ell)$ is an integer. (See Figure 5, where $k = 4$. It is true in general that the Bohr-Sommerfeld points are integer points of the moment map — see Proposition 9.1.) Since the image of μ_k will be the interval $[0, k]$, there will be $k + 1$ integer points in this interval, including the endpoints.

Cover \mathbb{R} by a sequence of pairwise overlapping intervals. This induces a covering of S^2 by sets U_j which are the equivalent in S^2 of bands in the model spaces. In fact, we have a symplectomorphism from such a set to a band in the cylinder, or a disc in \mathbb{C} . By the results in Section 8, the cohomology of U_j is isomorphic to the cohomology of the corresponding band or disc in the model space. This cohomology is trivial in all dimensions other than 1, as we have seen, and even in dimension 1 is only non-trivial if U_j contains a non-singular Bohr-Sommerfeld leaf. Adding up the results from all of the U_j by Mayer-Vietoris (Proposition 6.4 below), we have:

Theorem 5.1. The sheaf cohomology $H^q(S^2, \mathcal{J})$ is zero if $q \neq 1$, and has dimension equal to the number of non-singular Bohr-Sommerfeld leaves if $q = 1$.

Thus, when the prequantum line bundle is \mathbb{L}_k , the quantization has dimension $(k - 1)$.

6 The multidimensional case

For the case of higher dimensions, the model space we use is $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$, where k will be determined by the dimension of the leaf. In this section we describe the basic set-up of the model space, as well as prove some results about piecing together in the multidimensional case. Computation of the cohomology of the model space is postponed until the following section.

6.1 The model space

Let $M_0 = (\mathbb{R} \times S^1)^m \times \mathbb{C}^k$ (where for the moment we write m for $n - k$, simply for ease of notation), with coordinates

$$(t_1, \theta_1, \dots, t_m, \theta_m, s_1, \phi_1, \dots, s_k, \phi_k)$$

using the same conventions as in sections 3 and 4. In these coordinates, the standard symplectic form is given by

$$\omega = dt_1 \wedge d\theta_1 + \dots + dt_m \wedge d\theta_m + ds_1 \wedge d\phi_1 + \dots + ds_k \wedge d\phi_k,$$

which is equal to

$$d(t_1 d\theta_1 + \cdots + s_k d\phi_k). \quad (6.1)$$

The calculations, at least in the beginning, are exactly the same as in the one-dimensional case, except with more indices.

The polarization is $P = \text{span}\{\frac{\partial}{\partial\theta_i}, \frac{\partial}{\partial\phi_j}\}$ whose leaves are surfaces of constant t and s . If we map M_0 to \mathbb{R}^{m+k} by projecting onto the t and s coordinates, the polarization is given by the level sets of this map. Let \mathbb{L}_0 be the trivial bundle $M_0 \times \mathbb{C}$, which we make a prequantum line bundle by giving it a connection whose potential 1-form is $\Theta = t_1 d\theta_1 + \cdots + s_k d\phi_k$, so that $d\Theta = \omega$.

Definition 6.1. A *band* in $(\mathbb{R} \times S^1)^m$ is a set of the form $I \times (S^1)^m$ where I is an open rectangle, i.e. the product of intervals, in \mathbb{R}^m . More generally, a *band* in $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$ is a band in the preceding sense in $(\mathbb{R} \times S^1)^{n-k}$, times the product of discs centred at 0 in \mathbb{C}^k . Even more generally, a *band* in a manifold M is a set symplectomorphic to a band (in the preceding sense) in $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$.

6.2 The flat sections

Proposition 6.2. A section of \mathbb{L}_0 which is flat along the leaves locally has the form

$$\sigma = a(t_1, t_2, \dots, s_k) e^{i(t_1\theta_1 + t_2\theta_2 + \cdots + s_k\phi_k)}, \quad (6.2)$$

which we may write as

$$a(\mathbf{t}, \mathbf{s}) e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} e^{i\langle \mathbf{s}, \boldsymbol{\phi} \rangle}$$

where a is a smooth function.

Proof. Let $\sigma: U \rightarrow \mathbb{C}$ be a section of \mathbb{L}_0 , which we view as a \mathbb{C} -valued function using the canonical trivialization. Then σ is flat along the leaves if $\nabla_X \sigma = 0$ for all $X \in P$; this will be true iff

$$\begin{aligned} \nabla_{\frac{\partial}{\partial\theta_j}} \sigma &= 0 \\ \nabla_{\frac{\partial}{\partial\phi_l}} \sigma &= 0 \end{aligned} \quad (6.3)$$

for all appropriate values of j and l , since the connection is linear in X .

Applying the argument from the proof of Proposition 3.1 to each θ and ϕ component in turn (note that the connection potential 1-form has the same

form in each component here as it did in the proposition), we obtain the differential equations

$$0 = \frac{\partial\sigma}{\partial\theta_j} - \sigma it_j \quad 0 = \frac{\partial\sigma}{\partial\phi_j} - \sigma is_j \quad (6.4)$$

for all θ and ϕ coordinates. Together, these equations imply the desired result. \square

As before, we have

Proposition 6.3. The Bohr-Sommerfeld set of M_0 is $\mathbb{Z}^m \times \mathbb{N}^k \subset \mathbb{R}^{m+k}$.

Proof. The Bohr-Sommerfeld points are those with integer t values, and positive integral s values, by exactly the same argument as in the cases of the cylinder and the complex plane. \square

6.3 Multidimensional Mayer-Vietoris

Proposition 6.4. Let U and V be subsets of M which are each the union of leaves of the polarization. Then the sequence

$$\cdots \rightarrow H^1(U \cup V, \mathcal{J}) \rightarrow H^1(U, \mathcal{J}) \oplus H^1(V, \mathcal{J}) \rightarrow H^1(U \cap V, \mathcal{J}) \rightarrow H^2(U \cup V, \mathcal{J}) \rightarrow \cdots \quad (6.5)$$

of sheaf cohomology is exact.

Proof. Let \mathcal{A} be a cover of $U \cup V$, which induces a cover of U , V , and $U \cap V$ as in section 3.4. Consider the sequence

$$0 \rightarrow C_{\mathcal{A}}^*(U \cup V, \mathcal{J}) \xrightarrow{q} C_{\mathcal{A}}^*(U, \mathcal{J}) \oplus C_{\mathcal{A}}^*(V, \mathcal{J}) \xrightarrow{r} C_{\mathcal{A}}^*(U \cap V, \mathcal{J}) \rightarrow 0 \quad (6.6)$$

of Čech cochain complexes computed with respect to the cover \mathcal{A} . We claim this sequence is exact.

The argument to show this is very similar to the proof of Proposition 3.9, which is the same result in the case of the cylinder. Exactness at the left and the middle follows from properties of sheaves, exactly as in that proof.

Exactness at the right is less straightforward. To show exactness, we need surjectivity of r ; so let A be a set in the cover \mathcal{A} and h an element of $\mathcal{J}(A \cap U \cap V)$. We require a flat section f on $U \cap A$ and g on $V \cap A$ whose difference on $U \cap V \cap A$ is h .

Since U and V are unions of leaves, they can be written as $\pi^{-1}(B_U)$ and $\pi^{-1}(B_V)$ for some subsets B_U and B_V of B . Since B is a manifold with

corners, we can find a partition of unity on B_U, B_V . By composing with the projection $\pi: M \rightarrow B$, we obtain a partition of unity ρ_U, ρ_V for U and V . (We take $\text{supp}(\rho_U) \subset U$.)

Now let $f = \rho_V h$ and $g = -\rho_U h$. Then f is a section over U , which is flat along the leaves because h is flat along the leaves and ρ_V is constant on the leaves. Similarly, g is a leafwise flat section defined over V . It is clear that $f - g = h$ on $U \cap V \cap A$.

Therefore r is surjective, and the sequence (6.6) is exact.

The short exact sequence (6.6) gives a long exact sequence in cohomology, as usual, which is the sequence (6.4) with respect to the cover \mathcal{A} . Doing this for each \mathcal{A} yields such a sequence for every cover. By the same argument as in (3.10), the exactness passes to the limit, and we have (6.4). \square

Remark. The existence of the exact sequence (6.4) seems to be true for cohomology using arbitrary sheaves—see [Iv], III.5.10. However, it is instructive to see how the particular properties of our sheaf allow a direct proof of Mayer-Vietoris.

7 A better way to calculate cohomology

He that goeth forth and weepeth, bearing precious seed, shall doubtless come again with rejoicing, bringing his sheaves with him.

— Psalm 126:6

For the higher dimensional model spaces, the challenge with trying to compute the cohomology of a cover directly, as we did in sections 3 and 4, is that in order to adequately cover a set of higher dimensions, the covers, and thus the bookkeeping, become more and more complicated. Instead, we use a sheaf theoretic argument to obtain the cohomology. I am grateful to Ruxandra Moraru for suggesting this approach, and for explaining much of the sheaf theory to me.

The structure of this section is as follows. First, in 7.1, we discuss the theoretical tools we will use. Then, in section 7.2, we apply them to the 2-dimensional case (which we have already calculated), for practice. After outlining the upcoming calculations in 7.3, in section 7.4 and 7.5 we apply our tools to the higher-dimensional case.

7.1 Theory

7.1.1 Spectral sequences

We will not attempt to describe the theory of spectral sequences in detail here, but refer the reader to [BT]. We will briefly review some of the facts about spectral sequences that will be necessary for our calculations, which we will not attempt to state in full generality, but just enough to suffice for our purposes.

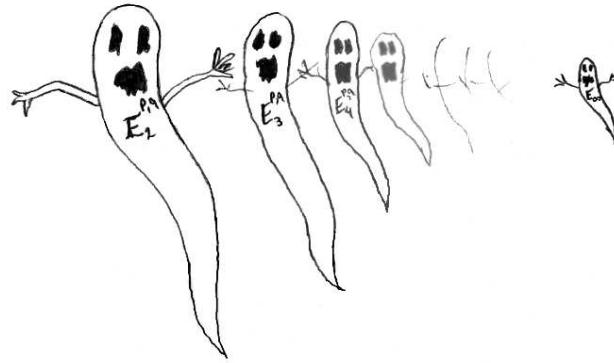


Figure 7: A spectral sequence

Recall that a *spectral sequence* is a collection $\{(E_r, d_r)\}$, where the E_r are vector spaces, the $d_r: E_r \rightarrow E_r$ are differentials (i.e. $d_r \circ d_r = 0$), and each group is the cohomology of the previous one, with respect to its differential: $E_{r+1} = H(E_r, d_r)$. Usually, E_r comes with a bigrading, and d_r shifts the bidegree, mapping $E_r^{p,q}$ to $E_r^{p+r, q-r+1}$.

A spectral sequence is often drawn in a chart, as in Figure 8, where each group $E^{p,q}$ is put in the appropriate square, and where we can think of each different r sitting on a different “page” in the diagram. The maps d_r go between the groups as shown; as r increases, the target of d_r for a fixed source moves down the diagonal.

If, for some r , all of the differentials are zero, then $E_{r+1} = E_r$, since $\ker d_r$ is everything and $\text{im } d_r$ is zero. If there is some s such that all d_r are zero for $r > s$, then all of the E_r 's for $r > s$ are the same, and we say the spectral sequence *stabilizes* or *converges*. We denote the common value of E_r for $r > s$ by E_∞ .

In many applications, a bigraded spectral sequence is used to obtain some singly-graded object. The grading on a spectral sequence is obtained from

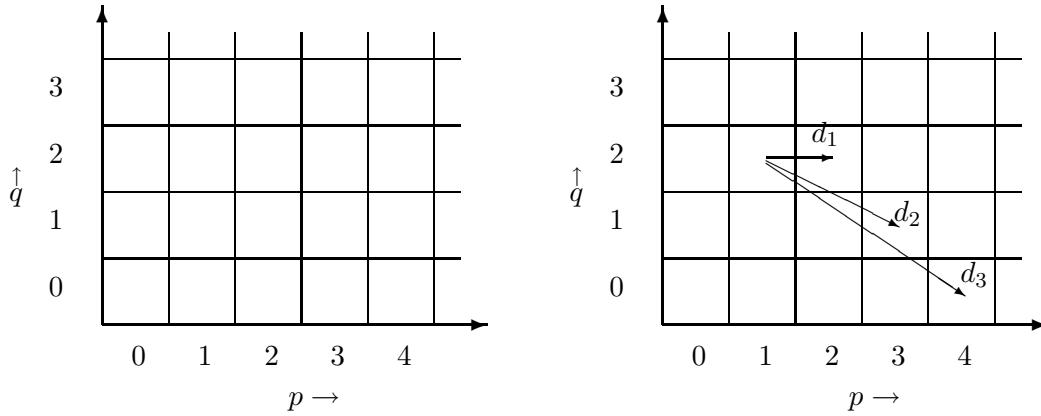


Figure 8: The diagram of a spectral sequence, and its differentials

the bigrading by summing along the diagonal:

$$E_r^k = \bigoplus_{p+q=k} E_r^{p,q}$$

If we say a spectral sequence converges to some singly-graded object, we mean it in this sense.

We observe that if, for some \$r \geq 2\$, the diagram of a spectral sequence has only one non-zero row, then the spectral sequence stabilizes at that value of \$r\$. More precisely:

Proposition 7.1. Let \$(E_r^{p,q}, d_r)\$ be a spectral sequence. Suppose that there is some number \$m\$ such that, for some \$s \geq 2\$,

$$E_s^{p,q} = 0 \quad \text{for all } q \neq m.$$

Then the spectral sequence stabilizes for \$r = s\$, i.e. \$E_\infty = E_s\$.

Proof. Since \$d_s\$ maps from \$E_s^{p,q}\$ to \$E_s^{p+s, q-s+1}\$, if \$s \geq 2\$, \$d_s\$ maps to a group with a different value of \$q\$. If there is only one value of \$q\$ for which the groups are non-zero, this implies that all the differentials must be zero. (See Figure 9, where all the blank boxes are zero, and a * represents something possibly nonzero.) Thus the spectral sequence stabilizes. \$\square\$

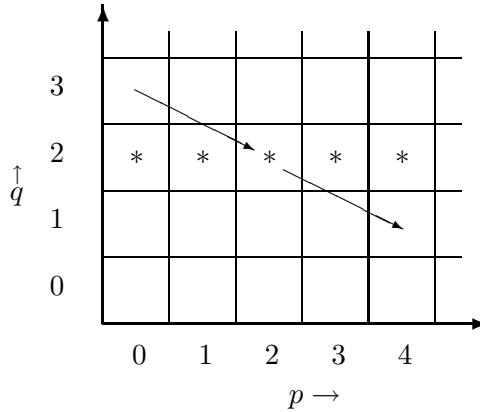


Figure 9: The differentials are all zero.

7.1.2 Leray spectral sequence

The spectral sequence we will be using is the following.

Theorem 7.2 ([G], Theorem 4.17.1). Let $f: M \rightarrow B$ be a continuous map, and let \mathcal{S} be a sheaf over M . Then there is a spectral sequence (called the *Leray spectral sequence*), whose E_2 term is given by

$$E_2^{p,q} = H^p(B; R^q f_* \mathcal{S}) \quad (7.1)$$

(where $R^q f_* \mathcal{S}$ are the direct image sheaves defined below) and which converges to $H^{p+q}(M, \mathcal{S})$.

Definition 7.3 ([GH], p. 463). Let $f: X \rightarrow Y$, and let \mathcal{S} be a sheaf on X . The q^{th} direct image sheaf $R^q f_* \mathcal{S}$ is the sheaf arising from the presheaf

$$U \mapsto H^q(f^{-1}(U), \mathcal{S}). \quad (7.2)$$

In the case $q = 0$, the sheaf is denoted simply by $f_* \mathcal{S}$ and called the *pushforward sheaf* of \mathcal{S} by f .

In the cases we consider, rather than find the higher direct image sheaves directly by using the above definition, we will find the *stalks* of the sheaves. The following result will be very useful in these calculations.

Theorem 7.4 (“Sheaf Theorist’s Trick”). Let $f: X \rightarrow Y$ be a proper map between locally compact spaces, and let \mathcal{S} be a sheaf on X . For $y \in Y$ and $q \in \mathbb{N}$ the restriction map

$$(R^q f_* \mathcal{S})_y \rightarrow H^q(f^{-1}(y); \mathcal{S}|_{f^{-1}(y)}) \quad (7.3)$$

(where $(R^q f_* \mathcal{S})_y$ denotes the stalk of the sheaf $R^q f_* \mathcal{S}$ at y , as in Definition 2.3) is an isomorphism.

Proof. This is Theorem III.6.2 from [Iv]. See also [G], Remarque 4.17.1. \square

In the cases we consider, the differentials will be trivial, and so usually

$$H^m(M, \mathcal{S}) = \bigoplus_{p+q=m} H^p(B, R^q \pi_* \mathcal{S}). \quad (7.4)$$

We will use the Leray spectral sequence *twice* in the course of this calculation: once with the map $(\mathbb{R} \times S^1)^n \rightarrow \mathbb{R}^n$, and once with $(S^1)^n \rightarrow (S^1)^{n-1}$.

Notation. We generally use \mathcal{J} to denote the sheaf of sections flat along the leaves. In what follows, use \mathcal{J}_n to denote this sheaf over $(\mathbb{R} \times S^1)^n$, when we need to be specific about the dimension. We will often need to consider one component of $(\mathbb{R} \times S^1)^n$ at a time, with other components fixed. We will use the notation \mathbf{t} to mean (t_1, \dots, t_n) as usual, but we will also use \mathbf{t}^{n-1} to mean (t_1, \dots, t_{n-1}) , to emphasize that we are not dealing with the coordinate t_n . We will use a similar notation $\boldsymbol{\theta}^{n-1}$.

7.1.3 Skyscraper sheaves

Definition. A *skyscraper sheaf* supported at a point p is a sheaf \mathcal{S} whose every stalk is zero except the stalk at p . More precisely, let A be some abelian group. Then $\mathcal{S}(U) = A$ if $p \in U$, and $\mathcal{S}(U) = 0$ otherwise. The group A is called the *tower* of the sheaf.

More generally, we allow a skyscraper sheaf to be supported at more than one point, provided the set of such points is discrete. Thus, for us, a skyscraper sheaf is one whose every stalk is zero outside of some discrete set.

Lemma 7.5 (Cohomology of a skyscraper sheaf). Suppose \mathcal{S} is a skyscraper sheaf supported on the discrete set $I \subset X$, with towers A_i respectively. Then the cohomology of \mathcal{S} is

$$\begin{aligned} H^0(X, \mathcal{S}) &= \bigoplus_{i \in I} A_i \\ H^q(X, \mathcal{S}) &= 0 \quad q > 0 \end{aligned} \quad (7.5)$$

Proof. This is a standard result. Since H^0 is just the global sections of the sheaf, the result for $q = 0$ is immediate from the definition; the result for $q > 0$ is, for example, Proposition IX.4.3 in [M]. \square

7.2 The case of one dimension

Consider the map $R \times S^1 \xrightarrow{\pi} \mathbb{R}$. Recall that elements of the sheaf \mathcal{J} locally have the form $a(t)e^{it\theta}$.

Let $I \subset \mathbb{R}$ be an open interval. By Definition 7.3, the pushforward sheaf $\pi_* \mathcal{J}$ is given by

$$\pi_* \mathcal{J}(I) \cong H^0(I \times S^1, \mathcal{J}).$$

This is just the set of global sections of \mathcal{J} over $I \times S^1$; as we argued in Section 3.2, this is 0 for all intervals I .

Similarly, the higher direct image sheaves are given by

$$(R^q \pi_* \mathcal{J})(I) = H^q(I \times S^1, \mathcal{J}).$$

The cohomology of $I \times S^1$ was computed in Section 3.2. From Proposition 3.28, we have that it is zero for all q , except when the interval I contains an integer, in which case $H^1 \cong \mathbb{C}^m$, where m is the number of integers in I (and all other H^q are still zero). Thus:

Lemma 7.6. The q th direct image sheaf $R^q \pi_* \mathcal{J}_1$ is 0 in the case $q \neq 1$, and a skyscraper sheaf supported on \mathbb{Z} , with all towers \mathbb{C} , if $q = 1$.

Proof. The case $q \neq 1$ is a restatement of the preceding paragraph. From the same discussion, supposing I is small,

$$R^1 \pi_* \mathcal{J}_1(I) = \begin{cases} \mathbb{C} & \text{if } I \text{ contains an integer} \\ 0 & \text{if } I \text{ contains no integer} \end{cases} \quad (7.6)$$

This is just the definition of a skyscraper sheaf. Its support is the set of integers. \square

7.3 The structure of the coming calculation

In the following section, we will calculate $H^m(U; \mathcal{J})$ for $U \subset (\mathbb{R} \times S^1)^n$. The calculation is somewhat complicated, and so we outline it here.

1. Let $\pi: (\mathbb{R} \times S^1)^n \rightarrow \mathbb{R}^n$ be the obvious projection, and let $\pi(U) = V$. The Leray spectral sequence gives $H^m(U; \mathcal{J})$ in terms of $H^p(V; R^q \pi_* \mathcal{J})$, so we need $R^q \pi_* \mathcal{J}$.

2. The sheaf theorist's trick gives a stalk of $R^q\pi_*\mathcal{J}$ as $H^q(\pi^{-1}(\mathbf{t}); \mathcal{J}')$, where $\mathbf{t} \in R^n$. Since $\pi^{-1}(\mathbf{t}) = \{\mathbf{t}\} \times T^n$, this amounts to finding $H^q(T^n; \mathcal{J}')$.
3. Let $\rho: T^n \rightarrow T^{n-1}$ be projection onto the first $(n-1)$ coordinates. The Leray spectral sequence applied to ρ gives us $H^m(T^n; \mathcal{J}')$ in terms of $H^p(T^{n-1}; R^q\rho_*\mathcal{J}')$.
4. The sheaf theorist's trick gives a stalk of $R^q\rho_*\mathcal{J}'$ as $H^q(\rho^{-1}(x); \mathcal{J}')$, where $x \in T^{n-1}$.
5. Finally, $\rho^{-1}(x) = \{x\} \times S^1$, and so the calculation of $H^q(\rho^{-1}(x))$ is reduced to the calculation of $H^q(S^1)$, which is essentially the same calculation as the one in section 3.2.

We fill in the details in the following section, in reverse order: The calculation in item 5 is carried out in Lemma 7.7, to find the sheaves in item 4. Item 3 is done in Lemma 7.8, where we use the Leray spectral sequence to find $H^m(T^n; \mathcal{J}')$ by an induction argument. We apply this to find a description of the sheaf $R^q\pi_*\mathcal{J}$ in Corollary 7.9. In Lemma 7.10, we apply the previous result to find $H^p(V; R^q\pi_*\mathcal{J})$, as in 1. And finally, in Theorem 7.11 we use this to find the sheaf cohomology of a band in $(\mathbb{R} \times S^1)^n$.

In section 7.5, we carry out the same calculation in the case where the leaf is partially singular. The idea is the same, but the calculation is simpler, and we only use the Leray spectral sequence once, applied to the map that projects out one of the singular components.

We remark that this outline shows the purpose of considering the stalks of the higher direct image sheaves, rather than trying to compute the sheaves themselves from Definition 7.3. If we were to follow through this calculation considering the sheaf over a small open set at each stage instead of the stalk, we would in step 5 be computing not $H^q(\{\text{pt}\} \times S^1; \mathcal{J}')$ but $H^q(W \times S^1; \mathcal{J})$, where W is some small neighbourhood in $T^{n-1} \times \mathbb{R}^n$. To do this directly by a method similar to section 3, we would have to use a cover of a $(2n-1)$ -dimensional set, which becomes combinatorially unmanageable. The sheaf theorist's trick permits us to use induction in step 5 instead.

7.4 The case of several dimensions: non-singular

As noted above, let \mathcal{J}_n be the sheaf of leafwise flat sections over $(\mathbb{R} \times S^1)^n$. Let \mathcal{J}'_n denote the restriction of \mathcal{J}_n to $\{\mathbf{t}\} \times (S^1)^n$. (The sheaf \mathcal{J}'_n will depend on the value of \mathbf{t} , though this is not made explicit in our notation.) Elements

of \mathcal{J}'_n will have the local form

$$a'(t_1, \dots, t_n) e^{i(t_1\theta_1 + \dots + t_n\theta_n)}, \quad (7.7)$$

where t_1, \dots, t_n are constant, and a' is the germ of a smooth function of the t variables. We will be very careful in the following to note what sheaf we are working with. Let

$$\pi: (\mathbb{R} \times S^1)^n \rightarrow \mathbb{R}^n$$

be the projection to the \mathbb{R} factors, and let

$$\rho: T^n \rightarrow T^{n-1}$$

be projection to the first $n - 1$ factors.

Lemma 7.7. The higher direct images of \mathcal{J}'_n are

$$R^q \rho_* \mathcal{J}'_n \cong \begin{cases} \mathcal{J}'_{n-1} & \text{if } t_n \in \mathbb{Z} \text{ and } q = 1 \\ 0 & \text{otherwise} \end{cases} \quad (7.8)$$

Proof. First, the case $q = 0$ is easy to see directly: Let U be a small open neighbourhood in T^{n-1} . The preimage $\pi^{-1}(U)$ is $U \times S^1$, and $(\rho_* \mathcal{J}'_n)(U)$ is just the set of elements of \mathcal{J}'_n over $U \times S^1$. As noted above in (7.7), elements of \mathcal{J}'_n can be written locally as $a' e^{i(t_1\theta_1 + \dots + t_{n-1}\theta_{n-1})} e^{it_n\theta_n}$, where $(\theta_1, \dots, \theta_{n-1}) \in U$ and θ_n ranges from 0 to 2π .

In order for such a creature to be defined on the whole set, it must have the same value for $\theta_n = 0$ as for $\theta_n = 2\pi$, *as germs of functions of the t variables*. This is impossible, even if $t_n \in \mathbb{Z}$ —for example, $e^{it\theta}$ and $e^{it\theta+2\pi it}$ are different as germs, even if their values are the same for $t \in \mathbb{Z}$. Thus there are no elements of \mathcal{J}'_n defined on all of $U \times S^1$, and so $\pi_* \mathcal{J}'_n = 0$.

Next, consider the case $q = 1$.

Recall that

$$(R^1 \rho_* \mathcal{J}'_n)_{\boldsymbol{\theta}} \cong H^1(\rho^{-1}(\boldsymbol{\theta}), \mathcal{J}'_n) = H^1(\boldsymbol{\theta} \times S^1, \mathcal{J}'_n) \quad (7.9)$$

View \mathcal{J}'_n as a sheaf over S^1 . This calculation is very similar to the one in section 3.2; the difference is that we will be dealing with germs in the t variables, instead of functions of t .

Cover S^1 with three sets E , F , and G , as in section 3.2. An element of \mathcal{J}'_n over S^1 looks like

$$a' e^{i(t_1\theta_1 + \dots + t_{n-1}\theta_{n-1})} e^{it_n\theta_n}$$

where $\theta_1, \dots, \theta_{n-1}$ are fixed, and θ_n ranges over all values from 0 to 2π , and all the t variables are fixed. This element is determined by a' , which is the germ of a smooth function of (t_1, \dots, t_n) .

A Čech 1-cochain is a collection of three elements of \mathcal{J}'_n , one for each intersection of E , F , and G . Thus such a cochain is determined by three germs $\{a'_{EF}, a'_{FG}, a'_{GE}\}$. All 1-cochains are cocycles.

The calculation of this cohomology is more or less the same as in section 3.2. In order for the cochain $\{a'_{EF}, a'_{FG}, a'_{GE}\}$ to be a coboundary, we need a 1-cochain $\{b_E, b_F, b_G\}$ which is a primitive.

This condition leads us to a set of equations like (3.7), except with some extra factors:

$$a_{EF}e^{it_n\theta_E}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} = b_F e^{it_n\theta_F}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} - b_E e^{it_n\theta_E}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} \quad \text{on } E \cap F \quad (7.10a)$$

$$a_{FG}e^{it_n\theta_F}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} = b_G e^{it_n\theta_G}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} - b_F e^{it_n\theta_F}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} \quad \text{on } F \cap G \quad (7.10b)$$

$$a_{GE}e^{it_n\theta_G}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} = b_E e^{it_n\theta_E}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} - b_G e^{it_n\theta_G}e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle} \quad \text{on } G \cap E \quad (7.10c)$$

Here each of the a 's and b 's are germs in t_1, \dots, t_n .

The argument following (3.7) in section 3.2 goes through unchanged, working with germs rather than functions. (The extra factor of $e^{i\langle \mathbf{t}, \boldsymbol{\theta} \rangle}$, since it is nonzero, can be cancelled from each equation. It is also worth noting that the role of t in section 3.2 is played here by t_n .) We get the same matrix (3.9) for the system, and the same condition that the system has a solution if $e^{-2\pi it_n} \neq 1$. We can think of t_n as a parameter of the sheaf, and if it is not an integer, then the cohomology in (7.9) is zero.

If $e^{-2\pi it_n} = 1$, then by the same linear algebra argument, the system has a solution only if

$$a_{EF}|_{t_n} + a_{FG}|_{t_n} + a_{GE}|_{t_n} = 0. \quad (7.11)$$

Thus the cohomology (7.9) is given by

$$H^1(\boldsymbol{\theta} \times S^1, \mathcal{J}'_n) \cong \{\text{germs in } t_1, \dots, t_n\} / \{a_{EF}|_{t_n} + a_{FG}|_{t_n} + a_{GE}|_{t_n} = 0\}. \quad (7.12)$$

Define a map from 1-cochains to \mathcal{J}'_{n-1} by

$$\{a'_{EF}, a'_{FG}, a'_{GE}\} \mapsto (a'_{EF} + a'_{FG} + a'_{GE})e^{i(t_1\theta_1 + \dots + t_{n-1}\theta_{n-1})}.$$

Its kernel is cocycles which have $(a'_{EF} + a'_{FG} + a'_{GE})|_{t_n} = 0$, i.e. coboundaries. Therefore,

$$\mathcal{J}'_{n-1}(U) \cong \{\text{cocycles}\} / \{\text{coboundaries}\} \cong H^1(U; \mathcal{J}'_n),$$

which equals $R^1\rho_*\mathcal{J}'_n$.

Finally, a similar technique works for the higher direct images—the argument of Section 3.2, with germs instead of functions, shows that $R^q\rho_*\mathcal{J}'_n = 0$ for $q \geq 2$. \square

Lemma 7.8. Assume $\mathbf{t} \in \mathbb{Z}^n$. Then

$$H^m(\{\mathbf{t}\} \times T^n; \mathcal{J}'_n) \cong \begin{cases} 0 & \text{if } m \neq n \\ \mathbb{C} & \text{if } m = n. \end{cases} \quad (7.13)$$

If $\mathbf{t} \notin \mathbb{Z}^n$, then $H^m(\{\mathbf{t}\} \times T^n; \mathcal{J}'_n) = 0$.

Proof. Assume first that all coordinates of \mathbf{t} are integers. We proceed by induction on n .

First, the case of $n = 1$ is just Lemma 7.6: $H^1 \cong \mathbb{C}$, $H^{\text{other}} = 0$.

Consider the Leray spectral sequence for the map ρ . It will have E_2 term given by

$$E_2^{p,q} = H^p(T^{n-1}; R^q\rho_*\mathcal{J}'_n).$$

By Lemma 7.7, all of the $R^q\rho_*\mathcal{J}'_n$ are zero except R^1 . Therefore, only one row of the spectral sequence is non-zero, and thus, as discussed in Proposition 7.1, the spectral sequence immediately stabilizes and $E_\infty = E_2$. Thus

$$\begin{aligned} H^m(T^n; \mathcal{J}'_n) &= E_\infty^m \cong \bigoplus_{p+q=m} H^p(T^{n-1}; R^q\rho_*\mathcal{J}_n) \\ &\cong H^{m-1}(T^{n-1}; R^1\rho_*\mathcal{J}'_n) \end{aligned} \quad (7.14)$$

where the last line is true because all R^q are zero except R^1 . Also by Lemma 7.7, $R^1\rho_*\mathcal{J}'_n \cong \mathcal{J}'_{n-1}$. Thus

$$H^m(T^n; \mathcal{J}'_n) \cong H^{m-1}(T^{n-1}; \mathcal{J}'_{n-1}).$$

The induction is complete.

Now, suppose that $\mathbf{t} \notin \mathbb{Z}^n$. Without loss of generality, suppose $t_n \notin \mathbb{Z}$. Then by Lemma 7.7, all the sheaves $R^q\rho_*\mathcal{J}'_n$ are zero, and so all summands in (7.14) are zero. \square

Corollary 7.9. The sheaf $R^q\pi_*\mathcal{J}_n$ over \mathbb{R}^n is 0 for $q \neq n$. When $q = n$, it is a skyscraper sheaf supported on \mathbb{Z}^n , with each tower isomorphic to \mathbb{C} .

Proof. Assume first that $\mathbf{t} \in \mathbb{Z}^n$.

By the sheaf theorist's trick, a stalk of $R^q\pi_*\mathcal{J}_n$ has the form

$$(R^q\pi_*\mathcal{J}_n)_{\mathbf{t}} \cong H^q(\pi^{-1}(\mathbf{t}), \mathcal{J}'_n) \cong H^q(T^n, \mathcal{J}'_n).$$

In Lemma 7.8, we just showed that this is isomorphic to \mathbb{C} , if $q = n$, and 0 otherwise. Therefore, the sheaf $R^1\pi_*\mathcal{J}_n$ has stalk \mathbb{C} when $\mathbf{t} \in \mathbb{Z}^n$, and zero otherwise. \square

Lemma 7.10. For $V \subset \mathbb{R}^n$,

$$H^p(V, R^q\pi_*\mathcal{J}) = \begin{cases} 0 & \text{if } q \neq n \\ \mathbb{C}^\nu & \text{if } p = 0 \text{ and } q = n \\ 0 & \text{if } p \geq 1 \end{cases} \quad (7.15)$$

Here ν is the number of points in the intersection $V \cap \mathbb{Z}^n$.

Proof. This all follows from Corollary 7.9.

The first case follows since $R^q\pi_*\mathcal{J} = 0$ if $q \neq n$, and so the cohomology is zero (the cohomology of a zero sheaf is zero).

The second case follows since the sheaf is a skyscraper. As given in Lemma 7.5, the 0^{th} cohomology of a skyscraper sheaf is the direct sum of the towers. Since $R^n\pi_*\mathcal{J}$ has tower \mathbb{C} over each point of \mathbb{Z}^n , the result for $p = 0, q = n$ follows.

The third case follows because the higher cohomology of a skyscraper sheaf is zero. \square

Theorem 7.11. For $U \subset (\mathbb{R} \times S^1)^n$ a band, $H^n(U, \mathcal{J}) \cong \mathbb{C}^\nu$, where ν is the number of Bohr-Sommerfeld orbits contained in U . For all other values of m , $H^m(U, \mathcal{J}) = 0$.

Proof. Let ν be the number of Bohr-Sommerfeld leaves contained in U , and let $V = \pi(U) \subset \mathbb{R}^n$. Since the Bohr-Sommerfeld leaves are precisely the fibres of π over the points in \mathbb{Z}^n , ν is the number of points in the intersection $V \cap \mathbb{Z}^n$.

At last, we will apply the Leray spectral sequence to the map $\pi: (\mathbb{R} \times S^1)^n \rightarrow \mathbb{R}^n$. The E_2 term of the sequence is

$$\mathcal{E}_2^{p,q} = H^p(V; R^q\pi_*\mathcal{J}). \quad (7.16)$$

According to Lemma 7.10, this is only non-zero if $(p, q) = (0, n)$. Thus the spectral sequence has only one non-zero entry (never mind one non-zero row), and so by Proposition 7.1, the spectral sequence stabilizes immediately and

$$H^m(U, \mathcal{J}) = \bigoplus_{p+q=m} H^p(V, R^q\pi_*\mathcal{J}). \quad (7.17)$$

The only non-zero summand is the $(0, n)$ one, which again by Lemma 7.10 is isomorphic to \mathbb{C}^ν . \square

7.5 The partially singular case

We now deal with the case where the Bohr-Sommerfeld leaf is singular. This means that the model space is $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$, for some $k > 0$. In this case, a simple neighbourhood of the leaf is a more general band: the product of bands around the cylinder components, and a product of discs in the \mathbb{C} components (see Definition 6.1).

Theorem 7.12. If $U \subset (\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$ is a band, and $k > 0$, then $H^m(U; \mathcal{J}) = 0$ for all values of m .

Proof. Assume, for definiteness, that $k = 1$, so $U \subset (\mathbb{R} \times S^1)^{n-1} \times \mathbb{C}$.

Let

$$p: U \rightarrow Y = (\mathbb{R} \times S^1)^{n-1}$$

be projection onto the first $n-1$ coordinates. The Leray spectral sequence for the map p gives $H^m(U; \mathcal{J})$ in terms of $H^r(Y; R^q p_* \mathcal{J}_n)$. By the sheaf theorist's trick (Theorem 7.4), the sheaf $H^r(Y; R^q p_* \mathcal{J}_n)$ is given by

$$(R^q p_* \mathcal{J})_x \cong H^q(p^{-1}(x); \mathcal{J}'_n), \quad (7.18)$$

where \mathcal{J}'_n is the sheaf \mathcal{J}_n restricted to $p^{-1}(x) = \{x\} \times D^2$.

Elements of the sheaf \mathcal{J} , recall, have the local form $a(t, s)e^{i\langle t, \theta \rangle}e^{i\langle s, \phi \rangle}$. Since we are assuming for the moment that $k = 1$, elements of \mathcal{J}' locally have the form $a'(\mathbf{t}, s)e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi}$, where a' is the germ of a smooth function of all the $(n-1)$ t variables, and a smooth function of s . In D^2 , the ϕ and s variables can change; the \mathbf{t} and θ variables are fixed by x (in fact $x = (\mathbf{t}, \theta)$). Call such a creature a "semigerm". Thus we wish to calculate

$$H^q(\{x\} \times D^2; \mathcal{J}'_n)$$

where \mathcal{J}'_n is the sheaf of semigerms.

This calculation is almost identical to the one in Section 4, except that the elements of the sheaf are of the form $a'(\mathbf{t}, s)e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi}$ instead of $a(s)e^{is\phi}$. The same calculation as the one following Proposition 4.4 goes through in this case, except that the coefficients h_{EF} , etc., are semigerms instead of just functions of s , and there is an extra factor of $e^{i\langle \mathbf{t}, \theta \rangle}$ multiplying everything. For example, equation (4.3) becomes

$$\begin{aligned} h_{FE}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_E} - h_{AF}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_F} + h_{AE}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_E} &= 0 && \text{on } AEF \\ h_{FG}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_F} - h_{AG}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_G} + h_{AF}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_F} &= 0 && \text{on } AFG \\ h_{GE}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_G} - h_{AE}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_E} + h_{AG}e^{i\langle \mathbf{t}, \theta \rangle}e^{is\phi_G} &= 0 && \text{on } AGE \end{aligned} \quad (7.19)$$

where each $h_{UV}(\mathbf{t}, s)$ is a germ in the t variables and a smooth function of s .

Regardless of the (fixed) values of \mathbf{t} and $\boldsymbol{\theta}$, $e^{\langle \mathbf{t}, \boldsymbol{\theta} \rangle}$ will be nonzero, and so it can be cancelled from each of the equations. Also, as in section 4.2, we can cancel the factors of $e^{is\phi}$, keeping track of the conventions on the ϕ coordinates (given in 4.2 just before equation (4.4)). This gives a set of equations identical to (4.4), this time for the semigerms h_{EF} , h_{FG} , etc., and the calculation proceeds exactly the same, and gives the same result for $\{x\} \times D^2$ as in Theorem 4.6, that the degree 1 cohomology is zero.

Similarly, the argument in the proof of Proposition 4.7 goes through unchanged, with semigerms and the extra factor of $e^{\langle \mathbf{t}, \boldsymbol{\theta} \rangle}$, and tells us that the cohomology in all other degrees is trivial.

Finally, we return to the calculation of $H^m(U; \mathcal{J})$. By what we have just shown, $R^q p_* \mathcal{J}$ is zero for all $q \geq 0$. Thus $H^r(Y; R^q p_* \mathcal{J}) = 0$ for any choice of r and q , and so the Leray spectral sequence for the map p not only stabilizes, it has all E_2 terms equal to zero. Therefore $H^m(U; \mathcal{J}) = 0$ for all m .

Note we have assumed for simplicity that $k = 1$, but the same argument will hold if $k > 1$ as well; there will just be more s and ϕ coordinates. Thus

$$H^m(U; \mathcal{J}) = 0 \quad \text{for all } m \tag{7.20}$$

whenever the leaf is at all singular. □

8 Piecing and glueing

The conditions in Definitions 2.8 guarantee that our spaces of interest have neighbourhoods which look like open sets in the model spaces. We will use the results on the cohomology of the model spaces to obtain results about the cohomology of these spaces. In this section we develop the theory necessary to transfer results about sheaf cohomology from one manifold to another.

The fact that our spaces are locally symplectomorphic to $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$ is a part of the definition. In section 8.2 we also obtain, with no extra hypotheses, a stronger condition on the symplectomorphism, which will enable us to compare sheaf cohomology. Finally, in 8.3, we put everything together to prove our main theorem about the cohomology of these spaces.

First, though, we need some theory.

8.1 Necessary sheaf theory

Lemma 8.1. Let \mathcal{S} and \mathcal{T} be sheaves over manifolds M and N , respectively, and $f: M \rightarrow N$ a smooth map. Suppose that, for every open set $V \subseteq N$, we

have a group homomorphism¹⁰

$$f^\sharp: \mathcal{T}(V) \rightarrow \mathcal{S}(f^{-1}(V))$$

which is compatible with both f and the restriction maps in the sheaves, in the sense that

$$(f^\sharp\alpha)|_{f^{-1}(V)} = f^\sharp(\alpha|_V) \quad \text{for all } V \text{ and } \alpha \in \mathcal{T}(V). \quad (8.1)$$

Then, for any open $V \subseteq N$, f^\sharp induces a map

$$f^*: H^*(f^{-1}(V), \mathcal{T}) \rightarrow H^*(V, \mathcal{S}).$$

Furthermore, this process is ‘functorial,’ in the sense that the composition of induced maps is the map induced by the composition.¹¹

Proof. In brief, the map f^\sharp induces a cochain map which is compatible with the coboundary, and thus induces a map on cohomology. The details are straightforward, and are left to the reader. \square

8.2 The induced map on cohomology

Theorem 8.2. Let N be a compact symplectic manifold with prequantization line bundle (\mathbb{L}_N, ∇^N) , equipped with a locally toric singular Lagrangian fibration. Let V be a neighbourhood of a leaf ℓ_N symplectomorphic to a neighbourhood U of a leaf ℓ_0 in a model space $M_0 = (\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$ (the existence of V and U is guaranteed by Definition 2.8). Then there exists an invertible map $f^\sharp: \mathcal{J}_{M_0}|_U \rightarrow \mathcal{J}_N|_V$.

The proof of Theorem 8.2 proceeds by a series of lemmas. Definition 2.8 guarantees the existence of a symplectomorphism $f: U \rightarrow V$. We first find a trivialization of \mathbb{L}_0 , with respect to which its connection has potential one-form $f^*\Theta_N$ (where Θ_N is the potential one-form of ∇^N). Then we use this trivialization to define f^\sharp in such a way that it takes flat sections to flat sections, proved in Lemma 8.6.

¹⁰In our applications, f^\sharp will be induced by f , hence the notation.

¹¹In more detail: Suppose we have maps

$$M \xrightarrow{f} N \xrightarrow{g} P,$$

and maps on sheaves

$$\mathcal{Q} \xrightarrow{g^\sharp} \mathcal{T} \xrightarrow{f^\sharp} \mathcal{S}.$$

Suppose further that $h = g \circ f$, $h^\sharp = f^\sharp \circ g^\sharp$, and all the requisite compatibilities are satisfied. Then $h^* = f^* g^*$ as maps on sheaf cohomology.

Lemma 8.3. We may choose ℓ_0 so that the holonomies of ℓ_0 and ℓ_N are equal.

Proof. As noted above, we already have a symplectomorphism $f: U \rightarrow V$. Changing t (the coordinate on \mathbb{R}^{n-k}) by a constant doesn't change the symplectic form on M_0 , and so we are free to choose t_0 , the value of t corresponding to $\ell_0 = f^{-1}(\ell_N)$.

Each leaf of the singular fibration is homeomorphic to T^{n-k} for some k —this is another consequence of the definition. Write $m = n - k$, as before. Let β_1, \dots, β_m be the set of fundamental cycles in T^m given by loops around each θ coordinate. These are then mapped by f to a set of fundamental cycles for ℓ_N , which we denote $\gamma_1, \dots, \gamma_m$.

The holonomy around γ_j is

$$\exp\left(i \int_{\gamma_j} \Theta_N\right)$$

(see section 2.1.1), while the holonomy around β_j is

$$\exp\left(i \int_{\beta_j} \Theta_0\right) = \exp\left(i \int_{\beta_j} \sum (t_a d\theta_a + s_b d\phi_b)\right) = e^{2\pi i t_j}.$$

Thus, if we choose the j^{th} coordinate of t_0 to be

$$\frac{1}{2\pi} \int_{\gamma_j} \Theta_N$$

for each $j = 1, \dots, m$, then the holonomies of ℓ_N and ℓ_0 will be equal. \square

Lemma 8.4. The bundle \mathbb{L}_N is trivializable over $V \subset N$.

Proof. The hypotheses on our spaces guarantee that we can choose V to be of the form $I^m \times T^m \times (D^2)^k$, with the leaf ℓ_N being identified with the central torus $\ell_N \cong \{\mathbf{t}^m\} \times T^m \times \{0\}$. A transverse neighbourhood is just $I^m \times D^{2k}$, which is contractible; therefore \mathbb{L}_N is trivializable over it. There is a free T^m action on V (just act on the T^m coordinate), which “sweeps out” the transverse disc over the neighbourhood V . This action gives us a trivialization over the whole neighbourhood. \square

Lemma 8.5. Choose a local trivialization of \mathbb{L}_N over V , and say Θ_N is the potential one-form of ∇^N with respect to this trivialization. Then there is a trivialization of \mathbb{L}_0 over U with respect to which ∇^0 has potential one-form $f^*\Theta_N$.

Proof. Let $\Theta_1 = f^*\Theta_N$ be the pullback of the potential one-form. We wish to show that there is a trivialization of \mathbb{L}_0 with respect to which the connection on M_0 has potential one-form Θ_1 .

Since the curvature of each connection is the symplectic form on the respective manifold, and f is a symplectomorphism, $d\Theta_0 = f^*d\Theta_N = d\Theta_1$. Therefore $\Theta_1 - \Theta_0$ is closed.

Recall from the proof of Lemma 8.3 that we have chosen t_0 so that

$$\int_{\gamma_j} \Theta_N = \int_{\beta_j} \Theta_0$$

for all j . Pushing the left side forward by f^{-1} , we obtain that

$$\int_{\beta_j} \Theta_1 = \int_{\beta_j} \Theta_0$$

for all j . Therefore $\Theta_1 - \Theta_0$ is closed and has integral 0 around all loops in T^k , and thus in ℓ_0 , and thus in U . Therefore, it is exact. (Note that the (s, ϕ) coordinates do not enter into this consideration, since U is a disc in those coordinates, and so there are no nontrivial loops in the (s, ϕ) coordinates.) Write $\Theta_1 = \Theta_0 + dG$ with $G: U \rightarrow \mathbb{R}$.

Now, suppose we change the trivialization of \mathbb{L}_0 by multiplying the fibres by some (nonzero) function $\psi: U \rightarrow S^1$; i.e., we take $(p, z) \mapsto (p, \psi(p)z)$. By Eq (2.4), this changes the potential one-form of a connection by subtracting $i\frac{1}{\psi} d\psi$ from it. Thus, we require a function ψ such that

$$i\frac{1}{\psi} d\psi = -dG.$$

Such a function is

$$\psi = e^{iG}.$$

Thus, multiplying the canonical trivialization of \mathbb{L}_0 by e^{iG} gives a new trivialization, with respect to which the connection ∇^0 has potential one-form Θ_1 . \square

Definition. Given N , M_0 , etc. as above, define the map $f^\sharp: \Gamma(V, \mathbb{L}_N) \rightarrow \Gamma(U, \mathbb{L}_M)$ as follows. Let s_V be the unit section of \mathbb{L}_N in some trivialization over V , and Θ_N the corresponding potential one-form of ∇^N . By Lemma 8.5 there is a trivializing section s_U of \mathbb{L}_0 over U , with respect to which ∇^0 has potential one-form $f^*\Theta_N$. If σ is a section of \mathbb{L}_N , then $\sigma = \phi s_V$ for a function ϕ . Then $f^\sharp \sigma$ is the section $(\phi \circ f) s_U$.

Note that f^\sharp is invertible, since f is invertible: map ϕs_U to $(\phi \circ f^{-1}) s_V$.

Lemma 8.6. If σ is flat along the leaves of N , then $f^\sharp\sigma$ is flat along the leaves of M_0 .

Proof. Since f is a diffeomorphism which carries leaves to leaves, it will suffice to prove that $\nabla_X^M f^\sharp \sigma = 0$ whenever $\nabla_{f_* X}^N \sigma = 0$. This follows from chasing the definitions, using a couple of facts about pullbacks and the chain rule. \square

Corollary 8.7. The map f^\sharp defined above is a sheaf map from $\mathcal{J}_{M_0}|_U$ to $\mathcal{J}_N|_V$.

Proof. Clearly f^\sharp is compatible with the restriction maps, and thus it is a sheaf map. Since it takes flat sections to flat sections, it maps \mathcal{J}_{M_0} to \mathcal{J}_N . \square

Given N , we have constructed an invertible map from \mathcal{J}_{M_0} to \mathcal{J}_N . Thus the proof of Theorem 8.2 is complete.

Corollary 8.8. Under the conditions of Theorem 8.2,

$$H^*(U, \mathcal{J}_M) \cong H^*(V, \mathcal{J}_N).$$

Proof. By Theorem 8.2, f^\sharp satisfies the conditions of Lemma 8.1, and thus induces a homomorphism f^* on cohomology. Note that f^\sharp is invertible, and its inverse g^\sharp satisfies the functoriality conditions in the footnote to Lemma 8.1. Therefore, f^* is an isomorphism, by the following standard argument: The maps induced on cohomology by f^\sharp and g^\sharp satisfy

$$f^* g^* = \text{id}; \quad g^* f^* = \text{id}$$

by functoriality, and so they must be isomorphisms. \square

8.3 Patching together

Finally, we are in a position to state and prove our main theorem. First a small lemma.

Lemma 8.9. Let M be a compact $2n$ -dimensional symplectic manifold, equipped with a locally toric singular Lagrangian fibration. Then M can be covered by finitely many open bands, such that any Bohr-Sommerfeld leaf is contained in only one band.

Proof. Recall first the definition of our spaces of interest, Definition 2.8.

Cover the base B by sets U which are homeomorphic to open rectangles in $\mathbb{R}^{n-k} \times \mathbb{R}_+^k$. By shrinking them, if necessary, we may assume that no

Bohr-Sommerfeld point lies in more than one of them. Then the inverse images of the U 's in M will be symplectomorphic to the inverse images of the rectangles in $\mathbb{R}^{n-k} \times \mathbb{R}_+^k$, which is precisely the definition of a band in this context (Definition 6.1), and no Bohr-Sommerfeld leaf will lie in more than one of them. Finally, finitely many of them will suffice to cover M since M is compact. \square

Theorem 8.10 (Main Theorem). Let M be a compact $2n$ -dimensional symplectic manifold, with a prequantization line bundle \mathbb{L} , and with a (singular) real polarization given by a locally toric singular Lagrangian fibration. Let \mathcal{J} be the sheaf of leafwise flat sections of \mathbb{L} . Then the cohomology groups $H^q(M; \mathcal{J})$ are zero for all $q \neq n$, and

$$H^n(M; \mathcal{J}) \cong \bigoplus_{b \in BS} \mathbb{C} \quad (8.2)$$

where the sum is taken over all non-singular Bohr-Sommerfeld fibres.

Proof. First, cover M by sets U as in Lemma 8.9. As each U is equivalent to a band in $(\mathbb{R} \times S^1)^{n-k} \times \mathbb{C}^k$, via a symplectomorphism satisfying the conditions of Theorem 8.2, it has the same cohomology as such a generalized band. By the results in section 7, a band has sheaf cohomology $H^q = 0$ for all $q \neq n$, and its n^{th} cohomology has one copy of \mathbb{C} for each non-singular Bohr-Sommerfeld leaf it contains. By Mayer-Vietoris (Prop 6.4), the cohomology of M is the sum of the cohomology of each of these generalized bands in the cover (since we are assuming no Bohr-Sommerfeld leaf lies in the intersection of any two sets in the cover), which gives us (8.2). \square

9 Real and Kähler polarizations compared

As noted in section 2.4, the (singular) foliation of a toric manifold by the fibres of the moment map is a singular Lagrangian fibration, and thus is a singular real polarization. Thus, by the above theorem, the quantization of a toric manifold has dimension equal to the number of nonsingular Bohr-Sommerfeld leaves.

Proposition 9.1. For M a toric manifold with moment map $\mu: M \rightarrow \Delta \subset \mathbb{R}^n$, the Bohr-Sommerfeld set is the set of integer lattice points in the moment polytope Δ . The singular Bohr-Sommerfeld points are the ones on the boundary of the polytope.

(This is a well-known result, but we include a proof here for the sake of completeness.)

Proof. Guillemin and Sternberg discuss the connection between Bohr-Sommerfeld points and action-angle variables in section 2 of [GS1]. Given a Lagrangian fibration $\pi: M \rightarrow B$ with compact fibres, they construct action coordinates as follows: Assume that $p \in B$ is a regular value of π . Choose a neighbourhood $V \subset B$ of p such that ω is exact on $\pi^{-1}(V)$, with β a primitive for ω . The fibres of π for points in V are tori; choose a homology basis $\gamma_1(q), \dots, \gamma_n(q)$ for the fibre over $q \in V$ which depends continuously on q . Define the action coordinates $a_j: V \rightarrow \mathbb{R}$ by

$$a_j(q) = \frac{1}{2\pi} \int_{\gamma_j(q)} \beta.$$

The $\{a_j\}$ are defined up to addition of a constant (which comes from changing β) and multiplication by an element of $GL(n, \mathbb{Z})$ (from changing the homology basis $\{\gamma_j\}$).

The holonomy around $\gamma_j(p)$ and $\gamma_j(q)$ differ by $\exp i(a_j(p) - a_j(q))$. Therefore, if the action coordinates are normalized so that $a_j(p) = 0$ at some Bohr-Sommerfeld point p , then another point q is a Bohr-Sommerfeld point iff all $a_j(q)$ are integers.

For a toric manifold, the coordinate system consisting of coordinates on $\Delta \subset \mathbb{R}^n$ together with coordinates on the torus fibre give action-angle coordinates, more or less by definition. Furthermore, because Δ is a polytope and thus simply-connected, these action-angle coordinates can be defined on all of the interior of the polytope. The discussion in [GS1] is for the case of a fibration, which here only applies to the regular values of μ , but it is not hard to extend the result to the singular values, as follows.

Assume that p_0 is a fixed point of the action; the case for a more general singular point is similar. By the Local Normal Form, we have a system of coordinates (s_j, ϕ_j) on a neighbourhood of the origin in \mathbb{C}^n , as in the local model space in section 6. Points with all s_j nonzero are regular points of μ , and the curves γ_j consisting of circles with fixed s_j in the (s_j, ϕ_j) plane form a homology basis for the regular fibres. Thus the s_j are action coordinates on the regular fibres, which extend continuously to the fixed point.

Assume that there is a Bohr-Sommerfeld point p in the neighbourhood on which these (s, ϕ) coordinates are defined. By the same argument as in section 4, all the $s_j(p)$ must be integers. The action coordinates t_j are also defined at p , since they are defined on the entire interior of Δ , and have integer

values at p . Therefore (as in (Eq 2.5) in [GS1]),

$$t_j = \sum a_{ij}(s_i + c_i)$$

for some $c \in \mathbb{Z}^n$ and $A = (a_{ij}) \in GL(n, \mathbb{Z})$. The above formula is valid when $s_j \neq 0$ for all j , but extends continuously to where $s_j = 0$, and so extends the definition of the action coordinates to all of M . Also, since $c \in \mathbb{Z}^n$ and $A = (a_{ij}) \in GL(n, \mathbb{Z})$, all the s_j are integers iff all the t_j are integers. Therefore the Bohr-Sommerfeld points in Δ are the integer lattice points.

Note that this gives another illustration of the fact (mentioned for example in Example 6.10 in [GGK]) that if M is quantizable, the fixed points are mapped to integer lattice points in \mathbb{R}^n .

Finally, as noted in section 2.3, if $x \in \Delta$ lies on a face of codimension k , then $\mu^{-1}(x)$ is an orbit of dimension $n - k$; thus, singular orbits correspond to points on the boundary of the polytope. \square

A toric manifold also has a natural complex structure, coming from its construction as a toric variety, and thus a natural Kähler polarization. If the manifold is quantized using this polarization, the dimension of the quantization is equal to the number of integer lattice points in the moment polytope, *including* points on the boundary. (This is a well-known result; see [H] for a more complete discussion.)

Thus the quantizations coming from Kähler and real polarizations are different, the difference being the number of lattice points lying on the boundary of the moment polytope (which is always non-zero, since fixed points map to lattice points).

References

- [A] J. E. Andersen, “Geometric Quantization of Symplectic Manifolds with Respect to Reducible Non-Negative Polarizations”, *Commun. Math. Phys.* **183**, 401-421 (1997)
- [BT] R. Bott and L. W. Tu, *Differential Forms in Algebraic Geometry*, Springer-Verlag, 1982
- [CdS1] A. Cannas da Silva, *Symplectic Toric Manifolds*, in M. Audin, A. Cannas da Silva, and E. Lerman, *Symplectic geometry of integrable Hamiltonian systems*, Advanced courses in mathematics, CRM Barcelona, Birkhäuser Verlag, Basel, 2003

- [CdS2] A. Cannas da Silva, *Lectures on Symplectic Geom.*, Springer, Berlin, 2001
- [D] Th. Delzant, “Hamiltoniens périodiques et images convexes de l’application moment,” *Bull. Soc. math. France*, **116** (1988), p 315–339
- [Dg] J. Dugundji, *Topology*, Allyn and Bacon, Boston, 1966
- [E1] L. H. Eliasson, *Normal forms for Hamiltonian systems with Poisson commuting integrals*, Ph.D. Thesis (1984).
- [E2] L.H. Eliasson, *Normal forms for Hamiltonian systems with Poisson commuting integrals—elliptic case*, Comment. Math. Helv. **65** (1990), no. 1, 4–35.
- [GGK] V. Guillemin, V. Ginzburg, and Y. Karshon, *Moment Maps, Cobordisms, and Hamiltonian Group Actions*, Mathematical Surveys and Monographs, **98**, American Mathematical Society, Providence, RI, 2002.
- [G] R. Godement, *Topologie Algébrique et Théorie des Faisceaux*, Hermann, Paris, 1964
- [GH] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*, John Wiley & Sons, New York, 1978
- [GS1] V. Guillemin and S. Sternberg, “The Gel’fand-Cetlin system and quantization of the complex flag manifolds”, *J. Funct. Anal.* **52** (1983), no. 1, 106–128
- [GS2] V. Guillemin and S. Sternberg, *Symplectic Techniques in Physics*, Cambridge University Press, 1984
- [H] M. D. Hamilton, “Quantization of a toric manifold is given by the integer lattice points,” in *Toric Topology*, Conference Proceedings of the International Conference on Toric Topology, Osaka, May 2006
- [Iv] B. Iversen, *Cohomology of Sheaves*, Springer-Verlag, Berlin, 1986
- [JW] L. Jeffrey and J. Weitsman, “Bohr-Sommerfeld Orbits in the Moduli Space of Flat Connections and the Verlinde Dimension Formula,” *Communications in Mathematical Physics*, **150** (1992), 593–630

- [Kg] M. Kogan, “On Completely Integrable Systems with Local Torus Actions,” *Annals of Global Analysis and Geometry*, **15** (1997), 543–553
- [K] B. Kostant, “Quantization and unitary representations. I. Prequantization,” *Lectures in modern analysis and applications, III*, Lecture Notes in Mathematics, **170**, pp. 87–208
- [MacL] S. MacLane, *Homology*, Springer-Verlag, Berlin, 1975
- [Mr] E. Miranda, *On symplectic linearization of singular Lagrangian foliations*, Ph.D. Thesis, Univ. de Barcelona, 2003
- [M] R. Miranda, *Algebraic Curves and Riemann Surfaces*, Graduate Studies in Mathematics, 5, American Mathematical Society, 1995
- [P] M. Puta, *Hamiltonian Mechanical Systems and Geometric Quantization*, Kluwer Academic Publishers, Dordrecht, 1993
- [S1] J. Śniatycki, “On Cohomology Groups Appearing in Geometric Quantization”, *Differential Geometric Methods in Mathematical Physics I* (1975)
- [S2] J. Śniatycki, *Geometric Quantization and Quantum Mechanics*, Applied Mathematical Sciences, **30**, Springer-Verlag, New York-Berlin, 1980
- [Th] W. P. Thurston, “Some Simple Examples of Symplectic Manifolds,” *Proc. Amer. Math. Soc.* **55** (1976), 467–468.
- [W] N.M.J. Woodhouse, *Geometric Quantization*, 2nd edition, Clarendon Press, Oxford, 1991
- [Z] Nguyen Tien Zung, “Symplectic topology of integrable Hamiltonian systems, I: Arnold-Liouville with singularities”, *Compositio Math.*, 101 (1996), 179–215